

IBM TotalStorage FAStT Storage Manager Version 8.4x



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

IBM TotalStorage FAStT Storage Manager Version 8.4x



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 137.

Second Edition (December 2003)

This edition replaces GC26-7622-00.

© **Copyright International Business Machines Corporation 2003. All rights reserved.**

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

Contents

Figures	vii
Tables	ix
About this document	xi
Who should read this document	xi
FAST installation process overview	xi
FAST documentation	xii
FAST Storage Manager Version 8.4 library	xiii
How this document is organized	xiii
Notices used in this document	xiv
Getting information, help and service.	xiv
Before you call	xiv
Using the documentation	xv
Storage Manager 8.4 online help	xv
Web sites.	xvi
How to send your comments.	xvi
Chapter 1. Introduction	1
Overview of Storage Manager 8.4	1
The storage management station	1
Storage Manager 8.4 online help.	2
Storage Manager 8.4 new features	2
Storage Manager 8.4 premium features	4
FAST600 restrictions	7
Overview of Storage Manager 8.4 software packages	7
SMruntime software package	7
SMclient software package	8
SMagent software package	8
SMutil software package	9
Failover driver.	9
Software installation sequence	9
Storage subsystem management methods.	10
Host-agent (in-band) management method.	10
Direct (out-of-band) management method	12
Installation process diagram	14
Operating system requirements	16
Setting up IP addresses for FAST storage controllers	17
Chapter 2. Installing storage management station software on AIX systems 19	
Hardware and firmware requirements	19
AIX restrictions	21
Removing previous versions of Storage Manager	22
Installing the client software	22
Prerequisites.	22
Installing SMruntime	23
Installing SMclient.	23
Installing the SMagent software (optional)	23
Prerequisites.	23
Installing SMagent	24
Starting or stopping SMagent	24
Configuring storage subsystems on AIX hosts	24
Adding storage subsystems to SMClient	24

Upgrading FAStT firmware and NVSRAM	25
Setting up an AIX host group.	26
Mapping LUNs to a FAStT storage partition	27
Creating direct and SAN-attached configurations	27
Creating a direct-attached configuration	28
Creating a SAN-attached configuration	28
Installing RDAC on AIX hosts	29
Prerequisites.	29
Installing the RDAC driver.	29
AIX configuration information.	30
Viewing and setting attributes of the RDAC driver for AIX	30
Viewing Object Data Manager attributes in AIX	35
Changing ODM attribute settings in AIX	36
Verifying the installation and configuration of AIX hosts	38
Identifying the controller ID numbers	39
Identifying device names and bus numbers	40
Identifying logical drives by operating system device names	46
Identifying FlashCopy volume types	48
Using fast I/O failure for fibre-channel devices	49
Using dynamic tracking of fibre-channel devices.	50
Using dynamic capacity expansion and dynamic volume expansion	51
Performing a dynamic capacity expansion operation	51
Performing a dynamic volume expansion operation	51
Resolving disk array errors	52
Redistributing volumes in case of failure	54

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

Hardware and firmware requirements	55
Installing the client software	56
Prerequisites.	56
Installing SMruntime	57
Installing the SMclient	57
Installing the host software	58
Prerequisites.	58
Installing the SMagent (optional)	58
Installing SMutil (required).	59
Configuring storage subsystems on HP-UX hosts	60
Adding storage subsystems to SMclient.	60
Upgrading FAStT firmware and NVSRAM	61
Setting up an HP-UX host group	62
Mapping LUNs to a FAStT storage partition	63
Creating direct-attached and SAN-attached configurations	63
Creating a direct-attached configuration	63
Creating a SAN-attached configuration	63
Enabling multipath I/O with PV-links	64
Using PV-links: Method 1	64
Using PV-links: Method 2	66
Redistributing volumes in case of failure	69

Chapter 4. Installing storage management station software on Solaris systems

Hardware and firmware requirements	71
Installing the client software	73
Prerequisites.	73
Installing SMruntime	73

Installing SMclient	74
Installing the host software	75
Prerequisites	75
Installing SMagent (optional)	76
Installing SMutil (required)	76
Configuring storage subsystems on Solaris hosts	77
Adding storage subsystems to SMclient	78
Upgrading NVSRAM and FAStT firmware	78
Creating direct-attached and SAN-attached configurations	79
Creating a direct-attached configuration	79
Creating a SAN-attached configuration	79
Configuring the host bus adapter cards	80
Installing the JNI adapter driver package	80
Modifying the JNI HBA settings	81
Installing the QLogic adapter driver package	86
Modifying the QLogic HBA settings	87
Setting up a Solaris host group	88
Mapping LUNs to a FAStT storage partition	89
Installing RDAC on Solaris hosts	89
Installing DMP on Solaris hosts	92
System requirements	92
DMP installation overview	92
Preparing for VERITAS DMP installation	92
Installing VERITAS Volume Manager packages for DMP	94
Installing the SMibmasl package	97
Configuring VERITAS Volume Manager	98
Starting the VERITAS Enterprise Administrator	101
Verifying external storage	102
Redistributing volumes in case of failure	103
Chapter 5. Completing the software configuration	105
Overview of heterogeneous host configurations	105
Configuring storage subsystems	105
Performing an initial automatic discovery of storage subsystems	105
Configuring alert notifications	107
Renaming storage subsystems	107
Creating a storage subsystem profile	107
Performing other storage subsystem management tasks	108
Chapter 6. Storage Manager 8.4 with high-availability cluster services	109
General information	109
Using cluster services on HP-UX systems	109
Using cluster services on Solaris systems	109
General hardware requirements	109
System dependencies	110
Using cluster services on AIX systems	111
General hardware requirements	111
High Availability Cluster Management Program	111
Parallel System Support Programs and General Parallel File System	112
GPFS, PSSP, and HACMP cluster configuration diagrams	112
Appendix A. AIX system requirements	119
Hardware requirements	119
Software requirements	119
RDAC installation requirements	119

Appendix B. HP-UX system requirements	121
Hardware requirements	121
Software requirements	121
Appendix C. Solaris system requirements	123
Hardware requirements	123
Software requirements	123
Appendix D. MC/Service Guard configuration details	125
Appendix E. JNI and QLogic host bus adapter settings	127
JNI HBA card settings	127
QLogic HBA card settings	129
Appendix F. Additional FAStT documentation	131
FAStT900 Fibre Channel Storage Server library	131
FAStT700 Fibre Channel Storage Server library	133
FAStT600 Fibre Channel Storage Server library	134
FAStT-related documents	135
Notices	137
Trademarks.	137
Glossary	139
Index	147

Figures

1. Installation process flow by current publications.	xii
2. Host-agent (in-band) managed storage subsystems	11
3. Direct (out-of-band) managed storage subsystems	13
4. Installation process (part 1)	14
5. Installation process (part 2)	15
6. Installation process (part 3)	16
7. Controller Properties window	40
8. Use of commands applied to validating system configuration.	41
9. Change Mappings window	48
10. Sample SMdevices command output (method 1)	65
11. Sample SMdevices command output (method 2)	66
12. Enterprise Management window.	106
13. Cluster configuration with single FAStT storage server — one to four partitions	113
14. Cluster configuration with three FAStT storage servers — one partition per FAStT	114
15. Cluster configuration with four FAStT storage servers — one partition per FAStT	115
16. RVSD cluster configuration with two FAStT storage servers — two partitions per FAStT	116
17. HACMP/GPFS cluster configuration with one FAStT storage server — one partition	117
18. HACMP/GPFS cluster configuration with two FAStT storage servers — two partitions per FAStT	118

Tables

1.	TotalStorage FASTT Storage Manager Version 8.4 titles by user tasks	xiii
2.	Maximum number of logical drives per partition	4
3.	Examples of possible configurations with more than 32 logical drives	4
4.	Using Storage Partitioning with FASTT storage servers	6
5.	Installation sequences of Storage Manager 8.4 software packages by host type	9
6.	FASTT storage servers that are supported on AIX systems	19
7.	Supported versions of hardware for AIX systems	20
8.	Attributes for dar devices	31
9.	Attributes for dac devices.	32
10.	Attributes for hdisk devices	33
11.	FASTT storage servers that are supported on HP-UX systems	55
12.	Supported versions of hardware for HP-UX systems	56
13.	Sample record of logical drive preferred and alternate paths	67
14.	FASTT storage servers that are supported on Solaris systems	71
15.	Supported versions of hardware for Solaris systems	72
16.	JNI model numbers, adapter drivers, and configuration files	80
17.	PSSP and GPFS software requirements	112
18.	Filesets required for AIX 4.3.3 RDAC	119
19.	Filesets required for AIX 5.1 RDAC.	119
20.	Filesets required for AIX 5.2 RDAC	120
21.	HP-UX kernel parameter configuration requirements	122
22.	Configuration settings for FCI-1063	127
23.	Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562	127
24.	Configuration settings for FCE-1063/FCE2-1063/FCE-6410/FCE2-6410	128
25.	Configuration settings for FC64-1063	129
26.	Configuration settings for QL2342	129
27.	TotalStorage FASTT900 Fibre Channel Storage Server document titles by user tasks	131
28.	TotalStorage FASTT700 Fibre Channel Storage Server document titles by user tasks	133
29.	TotalStorage FASTT600 Fibre Channel Storage Server document titles by user tasks	134
30.	TotalStorage FASTT related document titles by user tasks	135

About this document

Throughout this document, Storage Manager Version 8.4 refers to all version 8.4 software release levels. References to Storage Manager Version 8.41 refer exclusively to that release.

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 8.4 in AIX®, HP-UX, and Solaris.

Use this document to:

- Determine the hardware and software that is required to install the Storage Manager 8.4 software into your subsystem network
- Install the Storage Manager 8.4 software
- Upgrade the controller's nonvolatile storage random access memory (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 storage server. See Appendix F, "Additional FAStT documentation," on page 131 for a list of these documents.

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 the following concepts:

- Redundant array of independent disks (RAID)
- Small computer system interface (SCSI)
- Fibre channel technology

Readers should also have working knowledge of the applicable operating systems that are used with Storage Manager 8.4.

FAStT installation process overview

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

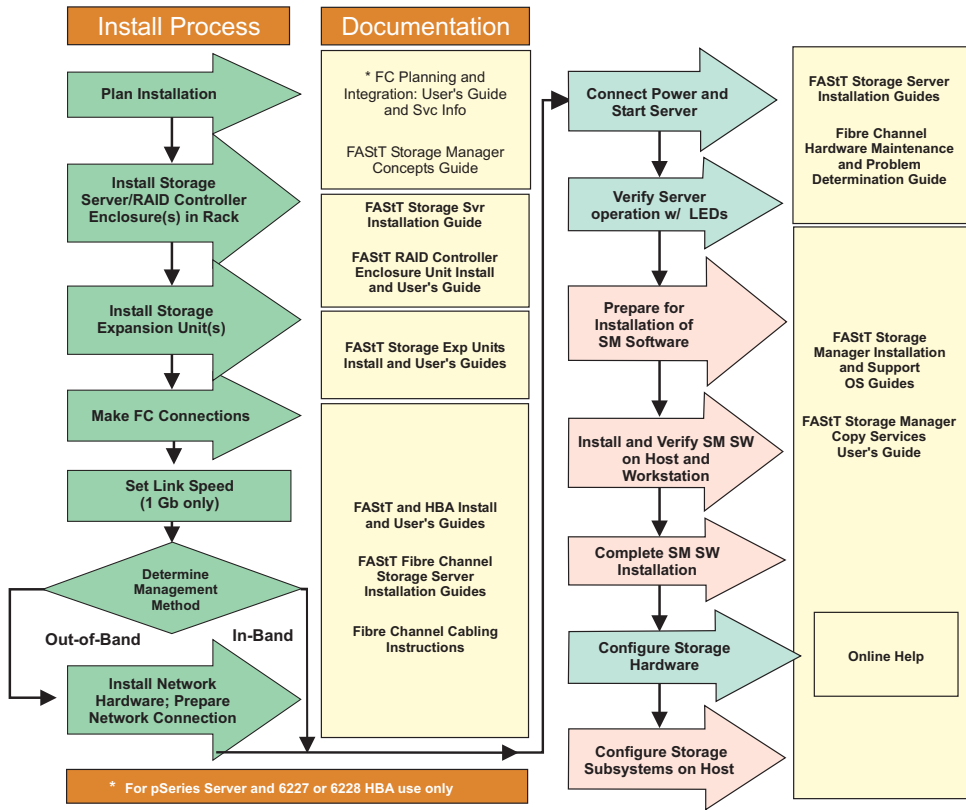


Figure 1. Installation process flow by current publications

FASTt documentation

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

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

FAStT Storage Manager Version 8.4 library

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

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

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
IBM TotalStorage FAStT Storage Manager 8.4x Installation and Support Guide for Intel-based Operating System Environments, GC26-7621	✓		✓	✓		
IBM TotalStorage FAStT Storage Manager 8.4x Installation and Support Guide for AIX, UNIX, and Solaris, GC26-7622	✓		✓	✓		
IBM TotalStorage FAStT Storage Manager Copy Services User's Guide, GC26-7561	✓		✓	✓	✓	
IBM FAStT Storage Manager Script Commands (see product CD)				✓		
IBM TotalStorage FAStT Storage Manager Concepts Guide, GC26-7560	✓	✓	✓	✓	✓	✓

How this document is organized

Chapter 1, "Introduction," on page 1 provides an introduction to the Storage Manager 8.4 product, including information about product resources in addition to this document.

Chapter 2, "Installing storage management station software on AIX systems," on page 19 provides information about how to install and use Storage Manager 8.4 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 55 provides information about how to install and use Storage Manager 8.4 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 71 provides information about how to install and use Storage Manager 8.4 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 configuration,” on page 105 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 8.4 with high-availability cluster services,” on page 109 provides information about high-availability clustering system options.

Appendix A, “AIX system requirements,” on page 119, Appendix B, “HP-UX system requirements,” on page 121, and Appendix C, “Solaris system requirements,” on page 123 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 8.4.

Appendix D, “MC/Service Guard configuration details,” on page 125 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 and QLogic host bus adapter settings,” on page 127 provides JNI host bus adapter (HBA) settings.

Appendix F, “Additional FASTT documentation,” on page 131 provides a reference list of other documentation in the FASTT600, FASTT700, and FASTT900 libraries, and other FASTT-related documents.

Notices used in this document

This document contains the following notices that are designed to highlight key information:

- Note:** These notices provide tips, guidance, or advice.
- Important:** These notices provide information or advice that might help you avoid inconvenient or problem situations.
- Attention:** These notices indicate possible damage to programs, devices or data. An attention notice is placed just before the instruction or situation in which damage could occur.

Getting information, help and service

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

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected.

- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system.
- Check for technical information, hints, tips, and new device drivers at the IBM support Web sites that are listed in this section.
- Use an IBM discussion forum on the IBM Web site to ask questions.

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

Using the documentation

Information about your IBM system and preinstalled software, if any, is available in the documents that come with your system. This includes printed books, online documents, readme files, and help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software.

Storage Manager 8.4 online help

Storage Manager 8.4 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 8.4. Click **Help** on the toolbar or press F1.

Note: The online help systems contain Storage Manager information that is common to all operating environments. For information about your particular operating system, see the following chapters and appendices of this document.

AIX operating systems

- Chapter 2, “Installing storage management station software on AIX systems,” on page 19
- Appendix A, “AIX system requirements,” on page 119

HP-UX operating systems

- Chapter 3, “Installing storage management station software on HP-UX systems,” on page 55
- Appendix B, “HP-UX system requirements,” on page 121

Solaris operating systems

- Chapter 4, “Installing storage management station software on Solaris systems,” on page 71
- Appendix C, “Solaris system requirements,” on page 123

Web sites

The most up-to-date information about FAStT storage servers and FAStT Storage Manager 8.4, including documentation and the most recent software, firmware, and NVSRAM downloads, can be found at the following Web sites.

FAStT Fibre Channel storage servers

Find the latest information about all of the FAStT 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 FAStT interoperability matrix

Find the latest information about operating system and HBA 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 8.4 readme files

Find the latest readme files for FAStT Storage Manager 8.4 at:

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 v8.4**. Accept the license agreement, then click the **v8.4** link for IBM FAStT Storage Manager for your operating system.

Fix delivery center for AIX

Find the latest AIX downloads at:

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

FAStT technical support

Find FAStT 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 FAStT600 or FAStT900 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/Building 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 8.4 in AIX, HP-UX, and Solaris environments.

Before you install Storage Manager 8.4, 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 v8.4**. Accept the license agreement, then click the **v8.4** 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 8.4 software. This document is available on the Storage Manager 8.4 installation CD.

Overview of Storage Manager 8.4

The IBM TotalStorage FASTT Storage Manager Version 8.4 is a set of client and host tools that enable you to manage IBM FASTT600, FASTT700, and FASTT900 Fibre Channel Storage Server subsystems from a storage management station.

The Storage Manager 8.4 software is available on the product CD. You can also download Storage Manager 8.4 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 v8.4**. Accept the license agreement, then click the **v8.4** 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 is also used to manage 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 8.4 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 8.4 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.

Storage Manager 8.4 new features

Storage Manager 8.4 provides support for the following new features:

- Support for QLogic host bus adapters on Solaris 8 operating systems
- Support for HACMP, PSSP, and GPFS cluster services
- Support for host-agent (in-band) management on AIX 5.1 and 5.2 operating systems
- Support for the FAStT EXP100 storage expansion unit on AIX 5.1 and 5.2 (see “Using the FAStT EXP100 expansion unit” for important information about this new feature)
- Support for up to 256 logical drives per storage partition, depending on the operating system (see “Performance planning” on page 3 for important information about this new feature)

Using the FAStT EXP100 expansion unit

The FAStT EXP100 storage expansion unit enables the attachment of Serial Advanced Technology Attachment (SATA) disk drives for FAStT600 and FAStT900.

Using the FAStT EXP100 requires that you upgrade the FAStT controller firmware to version 5.41.xx.xx. This firmware ships as part of the *FAStT Storage Manager 8.41 Support for FAStT EXP100 SATA Storage Expansion Unit CD*.

Notes:

1. Storage Manager 8.4 running firmware version 5.41.xx.xx is sometimes referred to as *Storage Manager 8.41*.
2. If you are not using the FAStT EXP100, do *not* upgrade the firmware to version 5.41.xx.xx.

3. The FAStT EXP100 is not supported on HP-UX or Solaris operating systems.

For more information about using the FAStT EXP100, be sure to read the *IBM TotalStorage FAStT EXP100 Storage Expansion Unit Release Notes* document, which ships with the FAStT EXP100.

Attention: It is essential that you review and understand the following issues *before* you upgrade the FAStT controller firmware to version 5.41.xx.xx to support the FAStT EXP100:

1. After you upgrade the FAStT storage server with controller firmware 5.41.xx.xx, you *cannot* downgrade the firmware to a prior level.
2. After you upgrade the FAStT storage server with controller firmware 5.41.xx.xx, you cannot attach FAStT600 and FAStT900 storage servers to fibre-channel expansion unit enclosures such as FAStT EXP500 or EXP700, even when these enclosures are not part of the same drive loop or loops as the FAStT EXP100s. You *must* disconnect any FAStT EXP500 or EXP700 enclosures from the FAStT storage server *before* you upgrade the controller firmware to version 5.41.xx.xx.

Note: IBM plans to enable the concurrent attachment of both the FAStT EXP100 SATA expansion unit and the FAStT EXP700 Fibre Channel expansion unit to a FAStT600 or FAStT900 storage server in future releases of the FAStT storage server firmware and FAStT Storage Manager host software.

3. After you install the FAStT controller firmware 5.41.xx.xx, any connected FAStT600 storage servers will *not* be able to use fibre-channel drives within the controller enclosure.

A FAStT600 storage server with controller firmware 5.41.xx.xx installed will not support any fibre-channel hard disk drives (FC HDDs) in the FAStT600 drive slots. You must remove any FC HDDs that are installed in these drive slots and replace them with drive blanks before upgrading the firmware.

4. Remote boot is not possible with FAStT storage servers that use the FAStT EXP100. You cannot use a logical drive that is created on the FAStT EXP100 as a boot disk or as a server operating system disk, due to the performance characteristics of SATA drive technology.

Performance planning

FAStT storage systems provide a high degree of flexibility to meet today's demanding storage applications. As such, care should be taken to plan configurations that meet specific application performance requirements.

Like most storage vendors, IBM publishes data sheets and specifications for the capabilities of the storage subsystem. These are intended for use as reference points for raw capabilities, but do not necessarily imply that actual performance specifications will be achieved for all configurations and operating environments. Many factors that are external to the FAStT storage server, such as server processor and memory, host bus adapters, command queue depth settings, and I/O size, can impact performance.

You must plan the configuration and setup of the FAStT for your operating environment. If proper planning is not performed, poor performance can result, such as low throughput or high disk latencies (response times).

With the release of Storage Manager 8.4, support of storage partitions with greater than 32 logical drives (up to 256) is another factor that must be taken into account during configuration planning. You can now connect many more logical drives to a

storage partition than was possible in Storage Manager 8.3. Depending on the I/O load requirements, it might be beneficial to spread volumes out on multiple partitions, if your operating system allows it.

Note: Not all operating system environments can support 256 logical drives per partition. See Table 2:

Table 2. Maximum number of logical drives per partition

Operating system	Maximum number of LUNs per partition
AIX	256
HP-UX 11.0	32
HP-UX 11.i	128
Solaris	256

The following example shows three possible configurations that have more than 32 logical drives.

Example: On previous versions of Storage Manager, connecting 64 logical disks to an AIX host required two storage partitions, each with 32 logical drives, and four HBAs. With Storage Manager 8.4, all 64 logical disks can be configured into one storage partition with 2 HBAs. You can also configure 256 logical disks into one partition with 2 HBAs. See Table 3:

Table 3. Examples of possible configurations with more than 32 logical drives

Configuration number	Storage Manager version	HBAs	Storage partitions	Logical drives
1	SM 8.3	4	2	64
2	SM 8.4	2	1	64
3	SM 8.4	2	1	256

You might expect the following performance results:

- Configuration 1 might perform better than configuration 2, which has more logical drives per storage partition.
- Depending on the application, configuration 3 might perform poorly because of the large number of logical drives per partition. Applications that do not require frequent use of (or connectivity to) many logical drives might respond quite well and meet performance expectations. Other applications, such as those that require simultaneous access to all volumes over extended periods of time, will likely experience lower throughput and high latencies on the individual logical drives.

For excellent information about how to plan your configuration, see the following documents:

- *IBM TotalStorage FAStT Storage Manager Concepts Guide*
- *IBM TotalStorage: FAStT Best Practices Guide*

Storage Manager 8.4 premium features

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

- Premium copy services features

- FlashCopy®
- VolumeCopy
- Remote Mirror Option
- Storage Partitioning premium feature (standard on all FAStT storage servers except the FAStT600)

FAStT600

The Remote Mirror Option is not currently supported on the FAStT600 or FAStT600 Turbo storage servers.

Note: Storage Partitioning is not a standard feature on the FAStT600.

Premium copy services features

This document lists the premium copy services features that are supported by Storage Manager 8.4, but does not describe them in detail. For detailed copy services concepts and procedures, see the following documents, which are included on the Storage Manager 8.4 installation CD-ROM.

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

Note: Premium copy services are not supported on Storage Manager 8.41.

Storage Manager 8.4 supports the following premium copy services:

FlashCopy FlashCopy supports the creation and management of FlashCopy logical drives, which are logical equivalents of complete physical copies. FlashCopy logical drives are created more quickly than physical copies, and they require less disk space.

FlashCopy is host addressable, so you can perform backups using FlashCopy while the base logical drive is online and user-accessible. When the backup completes, you can delete the FlashCopy logical drive or save it for reuse.

VolumeCopy VolumeCopy is used with FlashCopy to copy data from one logical drive to a target logical drive in a single storage subsystem. For example, you can use VolumeCopy for any of the following purposes:

- To copy data from arrays that use smaller capacity drives to arrays that use larger capacity drives
- To back up data
- To restore FlashCopy logical drive data to the base logical drive

Note: You must have FlashCopy enabled in order to use VolumeCopy. VolumeCopy can be purchased as part of a package with FlashCopy, or it can be purchased at a later time as an enhancement to FlashCopy.

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 I/O operations.

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 also defined in the storage partition.

Table 4 shows how many storage partitions are standard on your FAStT storage server.

Table 4. Using Storage Partitioning with FAStT storage servers

FAStT Storage Server	Storage Partitioning enabled by default	Maximum number of storage partitions
FAStT900	Yes (64 partitions standard)	64
FAStT700	Yes (64 partitions standard)	64
FAStT600	No	16
FAStT600 with Turbo option	Yes (8 partitions standard)	64

If you do not use Storage Partitioning, all logical drives are contained in a default host group. The 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 the systems that have the Storage Partitioning feature enabled, 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.

FAST600 restrictions

If you are using a FAST600 storage server, note the following restrictions:

- Without the Turbo option, the FAST600 supports only 56 drives and 16 storage partitions. Also, it does not support the 1 Gb host interface.
- The FAST600 and FAST600 Turbo do not support FAST EXP500 Fibre Channel expansion units.

Overview of Storage Manager 8.4 software packages

Storage Manager 8.4 contains the following client software packages:

SMruntime Storage Manager 8 Java™ compiler

SMclient Storage Manager 8 client package

Storage Manager 8.4 contains the following host software packages:

SMagent Storage Manager 8 agent package

SMutil Storage Manager 8 utility package

RDAC Storage Manager 8 multipath device drivers

Note: In addition to this document, see the *IBM FAST Storage Manager Concepts Guide* for more information about Storage Manager 8.4 software.

SMruntime software package

The Storage Manager 8.4 runtime software, SMruntime, is a Java compiler for the Storage Manager 8.4 client software, SMclient. SMruntime must be installed before SMclient is installed.

SMclient software package

The Storage Manager 8.4 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 FAStT600, FAStT700 and FAStT900 Fibre Channel Storage Servers
- FAStT EXP500 and EXP700 Fibre Channel expansion units
- FAStT EXP100 SATA expansion units

Note: The EXP100 is supported only on FAStT600 and FAStT900 storage servers running AIX 5.1 or 5.2.

Specifically, SMclient enables 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 partitions 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 NVSRAM
- View the major events log (MEL)
- Assign redundant RAID controllers
- Expand storage capacity using any of the following expansion units:
 - FAStT EXP500 Fibre Channel expansion unit
 - FAStT EXP700 Fibre Channel expansion unit
 - FAStT EXP100 SATA expansion unit

Note: The EXP100 is supported only on FAStT600 and FAStT900 storage servers running AIX 5.1 or 5.2.

The SMclient contains two main components:

Enterprise Management

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

Subsystem Management

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

SMagent software package

The Storage Manager 8.4 software agent package, SMagent, contains 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.

SMagent 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 SMagent, see “Host-agent (in-band) management method” on page 10.

SMutil software package

You can use the Storage Manager 8.4 utility, SMutil, to register and map new logical drives to the operating system.

Install SMutil on all HP-UX and Solaris host systems that are attached to a storage subsystem. The host computers are attached to the storage subsystem through the fibre channel.

Failover driver

AIX and Solaris host systems require one of the following failover drivers for fibre-channel path redundancy:

AIX	Redundant Disk Array Controller (RDAC)
Solaris	RDAC, or VERITAS Volume Manager with Dynamic Multipathing (DMP)

RDAC monitors I/O paths. If a component failure occurs in one of the fibre-channel paths, the RDAC reroutes all I/O to another path.

Note: The AIX RDAC driver files are not included on the Storage Manager 8.4 installation CD. You must follow the instructions in “Installing RDAC on AIX hosts” on page 29 to download them from the appropriate Web site.

Software installation sequence

Install the Storage Manager 8.4 software packages in the sequences shown in Table 5.

Table 5. Installation sequences of Storage Manager 8.4 software packages by host type

Step	AIX	HP-UX	Solaris
1	SMruntime (see “Installing SMruntime” on page 23)	SMruntime (see “Installing the client software” on page 56)	SMruntime (see “Installing SMruntime” on page 73)
2	SMclient (see “Installing SMclient” on page 23)	SMclient (see “Installing the client software” on page 56)	SMclient (see “Installing SMclient” on page 74)
3	SMagent (see “Installing the SMagent software (optional)” on page 23)	SMagent (see “Installing the SMagent (optional)” on page 58)	SMagent (see “Installing SMagent (optional)” on page 76)
4	RDAC (see “Installing RDAC on AIX hosts” on page 29)	SMutil (see “Installing SMutil (required)” on page 59)	SMutil (see “Installing SMutil (required)” on page 76)
5			RDAC or DMP (see “Installing RDAC on Solaris hosts” on page 89 or “Installing DMP on Solaris hosts” on page 92)

Storage subsystem management methods

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

Host-agent (in-band) management method

Using this method, you manage the storage subsystems through the fibre-channel I/O path to the host.

Direct (out-of-band) management method

Using this method, you manage the storage subsystems directly over the network through the Ethernet connection to each controller.

The management methods are described in the next sections.

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 11 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.

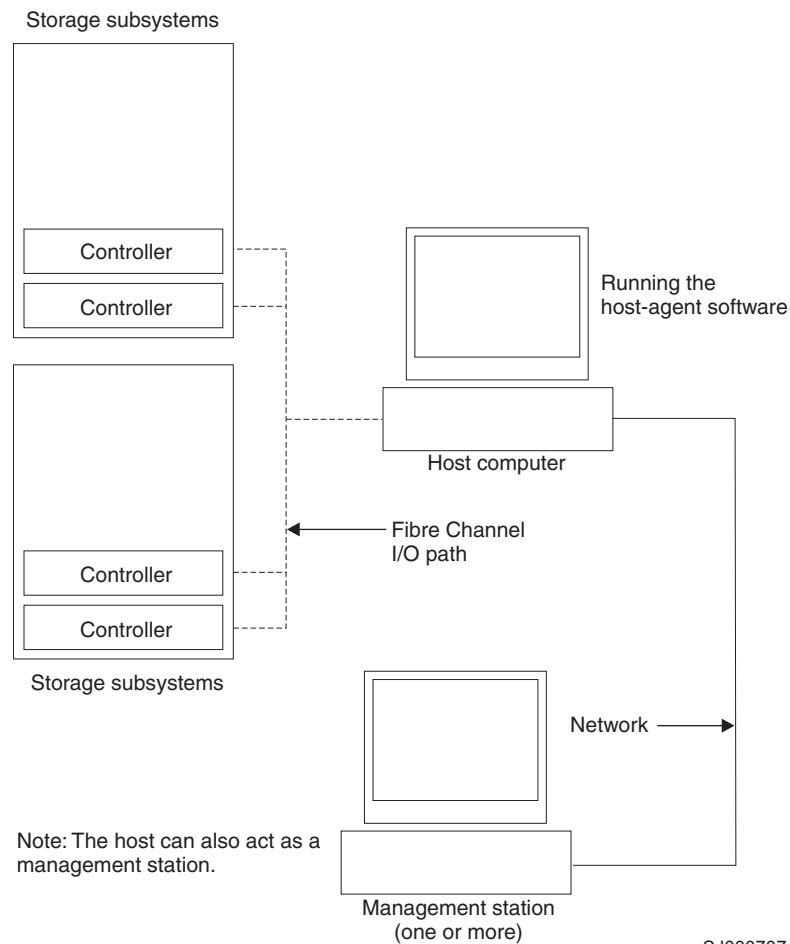


Figure 2. Host-agent (in-band) managed storage subsystems

Host-agent (in-band) management method advantages

Using the host-agent (in-band) management method provides 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 19, Chapter 3, “Installing storage management station software on HP-UX systems,” on page 55, or Chapter 4, “Installing storage management station software on Solaris systems,” on page 71).
- 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

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 13 shows the direct (out-of-band) management method.



SJ000708

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

If you use the direct (out-of-band) management method, see “Setting up IP addresses for FAST storage controllers” on page 17.

Direct (out-of-band) management method advantages

Using the direct (out-of-band) management method provides 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) management method disadvantages

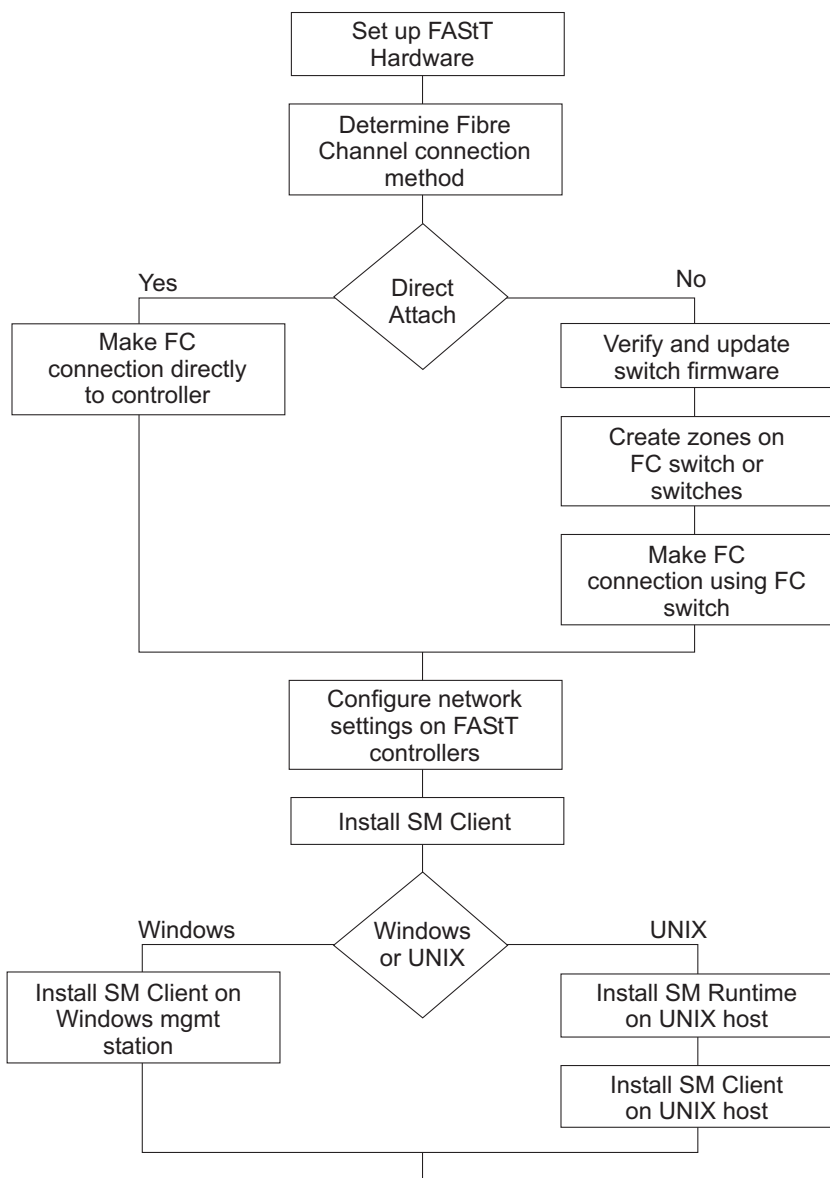
Using the direct (out-of-band) management method has the following disadvantages:

- 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 15, and Figure 6 on page 16 provide an overview of the installation process for out-of-band management.

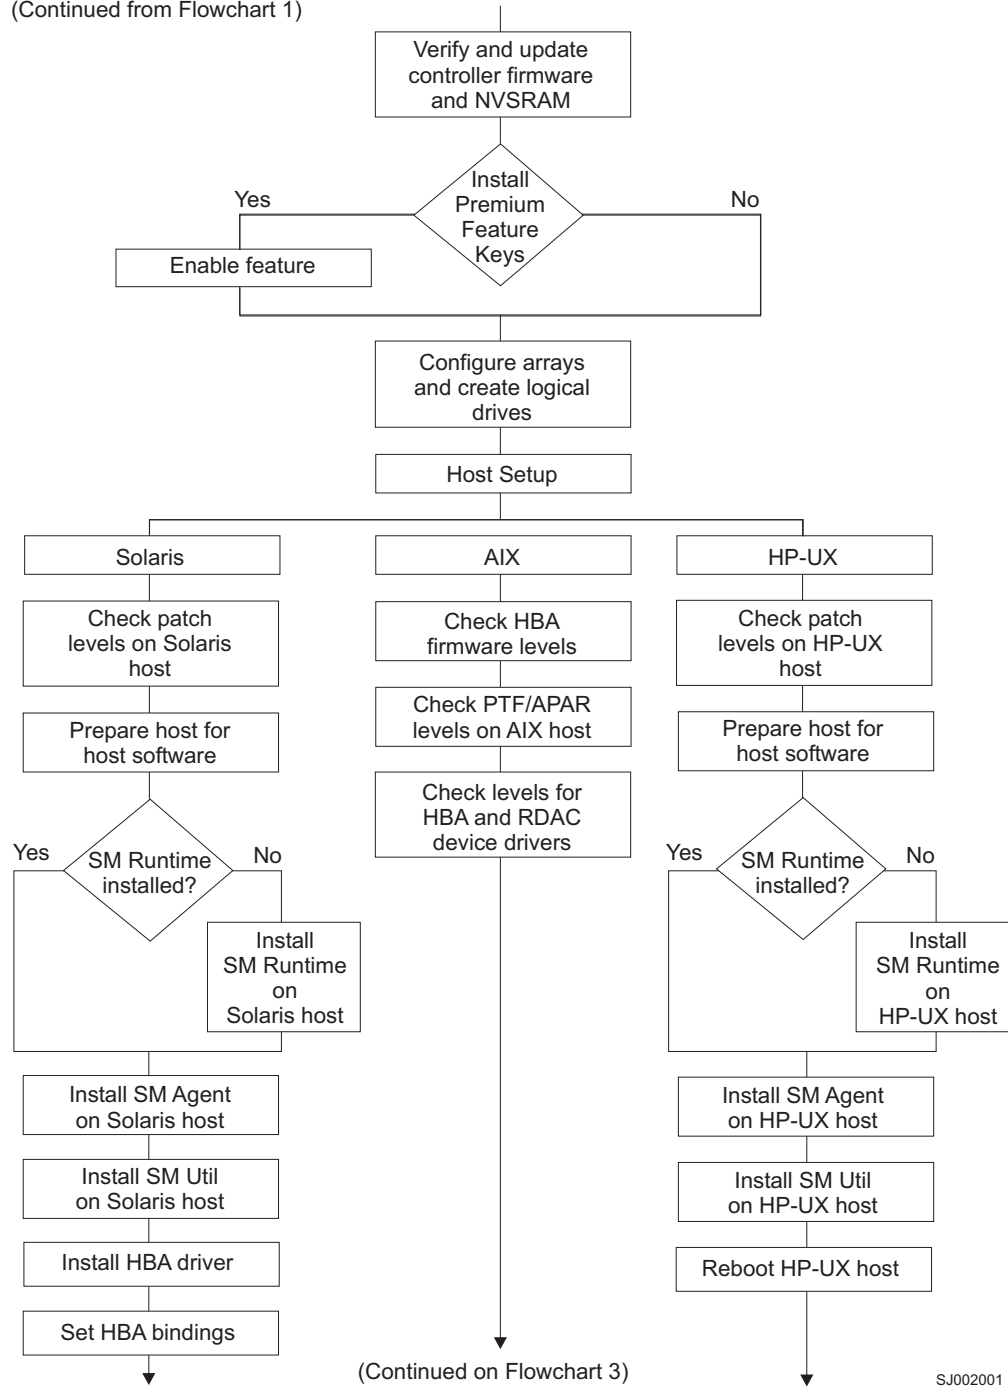


(Continued on Flowchart 2)

SJ002000

Figure 4. Installation process (part 1)

(Continued from Flowchart 1)

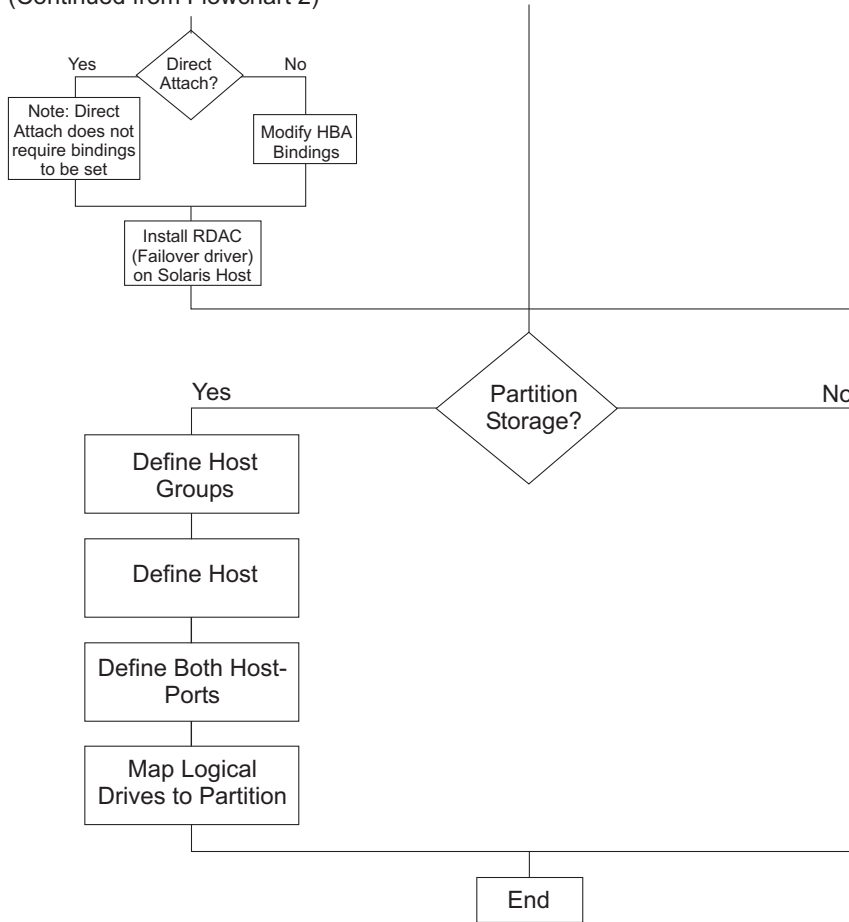


(Continued on Flowchart 3)

SJ002001

Figure 5. Installation process (part 2)

(Continued from Flowchart 2)



SJ002002

Figure 6. Installation process (part 3)

Operating system requirements

Storage Manager 8.4 supports the following IBM FAStT Fibre Channel Storage Servers in the following environments.

Operating systems	FAStT storage servers	Where to find more information
IBM AIX versions: 4.3.3 5.1 5.2	FAStT600, FAStT600 Turbo, FAStT700, FAStT900	Appendix A, "AIX system requirements," on page 119
HP-UX versions: 11.0 (32 bit or 64 bit) 11.11 11i	FAStT600, FAStT600 Turbo, FAStT700, FAStT900	Appendix B, "HP-UX system requirements," on page 121

Operating systems	FAStT storage servers	Where to find more information
Sun Solaris versions: 7 8	FAStT600, FAStT600 Turbo, FAStT700, FAStT900	Appendix C, "Solaris system requirements," on page 123

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:

Bits/second 57600

Data bits 8

Parity None

Stop Bit 1

Flow Control None

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 on page 17 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 information for AIX operating systems:

- “Hardware and firmware requirements”
- “AIX restrictions” on page 21
- “Removing previous versions of Storage Manager” on page 22
- “Installing the client software” on page 22
- “Installing the SMagent software (optional)” on page 23
- “Configuring storage subsystems on AIX hosts” on page 24
- “Creating direct and SAN-attached configurations” on page 27
- “Installing RDAC on AIX hosts” on page 29
- “AIX configuration information” on page 30
- “Using fast I/O failure for fibre-channel devices” on page 49
- “Using dynamic tracking of fibre-channel devices” on page 50
- “Using dynamic capacity expansion and dynamic volume expansion” on page 51
- “Resolving disk array errors” on page 52
- “Redistributing volumes in case of failure” on page 54

Attention: For cluster configurations, complete all applicable configuration procedures for each storage subsystem *before* installing the Storage Manager 8.4 software on a second host or cluster server.

Hardware and firmware requirements

Table 6 lists the FAStT Fibre Channel storage servers that are supported with Storage Manager 8.4.

Table 6. FAStT storage servers that are supported on AIX systems

Product Name	Model	Minimum appware/snapware (firmware) and NVSRAM versions
IBM TotalStorage FAStT900 Fibre Channel Storage Server	1742-90U	Storage Manager 8.40: Appware 5.40.06.xx NVSRAM N1742F900R840V02 Storage Manager 8.41 with EXP100: Appware 5.41.10.xx NVSRAM N1742F900R841V02
IBM TotalStorage FAStT700 Fibre Channel Storage Server	1742-1RU	Appware 5.40.06.xx NVSRAM N1742F700R840V03
IBM TotalStorage FAStT600 and FAStT600 Turbo Fibre Channel Storage Servers	1722-60U, 1722-60X	Storage Manager 8.40: Snapware 5.40.07.xx NVSRAM N1722F600R840V04 Storage Manager 8.41 with EXP100: Snapware 5.41.10.xx NVSRAM N1722F600R840V01

Note: Firmware version 5.41.10.xx supports only FAStT EXP100 SATA expansion units. Firmware version 5.40.06.xx supports only FAStT EXP500 and EXP700 storage expansion units. For more information, see “Storage Manager 8.4 new features” on page 2.

Table 7 lists the supported versions of other hardware to use with Storage Manager 8.4.

Table 7. Supported versions of hardware for AIX systems

Product Name	Model	Product release and firmware version
IBM TotalStorage FAStT EXP100 SATA Storage Expansion Unit	1710-10U or 10X	For firmware version, see the readme file that is shipped with this product. The readme file is also available at the Web site that is listed at the end of this table.
IBM TotalStorage FAStT EXP700 Fibre Channel Storage Expansion Unit	1740-1RU	ESM 9325
IBM TotalStorage FAStT EXP500 Fibre Channel 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 or higher
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 v8.4**. Accept the license agreement, then click the **v8.4** link for IBM FAStT Storage Manager for AIX.

AIX restrictions

The following restrictions apply to FASSt600, FASSt700, and FASSt900 storage servers:

- AIX 4.3.3 does not support greater than 32 LUNs.
- F-RAID Manager is not supported.
- The maximum number of partitions per AIX host, per FASSt storage server, is two.
- If you create more than 32 LUNs on a partition, you cannot use the release CD to install AIX on a FASSt device on that partition.

Important: If your system is booted from a FASSt device, do not create more than 32 LUNs on the partition that you are booting from.

Support may become available in a future release of AIX.

- AIX hosts (or LPARs) can support multiple HBAs and FASSt devices. However, there is a restriction on the number of HBAs that can be connected to a single FASSt storage server. You can configure up to two HBAs per partition and up to two partitions per FASSt storage server. Additional HBAs can be added to support additional FASSt 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 FASSt minihub is allowed.

Note: FASSt600 storage servers do not have minihubs.

- Single HBA configurations are allowed, but each single HBA configuration requires that both controllers in the FASSt 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 FASSt 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 “Setting up an AIX host group” on page 26.

Exception: If you are running a FASSt600 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 from a FASSt device is supported only on AIX 5.1 and 5.2.
- You cannot boot from a FASSt device that has one or more FASSt EXP100 SATA expansion units attached.
- When you boot from a FASSt device, both paths to the FASSt storage server 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.
- SMagent and in-band management are not supported on Storage Manager 8.41 .

Removing previous versions of Storage Manager

If you are upgrading from Storage Manager 8.xx, you do not have to remove the previous version. Installing Storage Manager 8.4 overwrites previous 8.xx versions.

If you are upgrading from Storage Manager 7.xx, complete the following steps to remove it before you install Storage Manager 8.4:

1. List all the Storage Manager software that is installed on the current AIX system, by typing the following command.

```
ls1pp -L | grep SM
```

2. Remove the Storage Manager software, by typing the following command.

```
installp -u SM_component_name
```

where *SM_component_name* is the name of a Storage Manager package component that you want to remove.

It might take several minutes for the task to complete. The system might be unresponsive during this time.

3. Repeat step 2 to remove each Storage Manager software package.

After you have completed these steps, you have finished removing the previous version of Storage Manager. You can now install Storage Manager 8.4 software.

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 19 and “Software requirements” on page 119.
- The correct file sets are present on the system.

Note: If the file sets are not present, follow the instructions in “Installing RDAC on AIX hosts” on page 29 to download them from the appropriate Web site.

- For a list of AIX 4.3.3 file sets, see Table 18 on page 119.
- For a list of AIX 5.1 file sets, see Table 19 on page 119.
- For a list of AIX 5.2 file sets, see Table 20 on page 120.

Installing SMruntime

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

1. Install SMruntime by typing the following command:

```
# installp -a -d /path_name  
/SMruntime.aix-08.40.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:

```
# lslpp -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-08.41.G5.01.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:

```
# lslpp -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.

Installing the SMagent software (optional)

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

Note: SMagent and in-band management are not supported on Storage Manager 8.41.

Prerequisites

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

- This is the AIX server that you identified as the storage management station.
- This machine meets the minimum hardware and software requirements described in Appendix A, "AIX system requirements," on page 119.
- SMruntime has been installed.

Installing SMagent

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

Note: The installation process adds an entry to the `/etc/inittab` file that starts SMagent each time the system is booted.

1. Install SMagent by typing the following command:

```
# installp -a -d /path_name/  
SMagent.aix-08.40.G5.00.bff SMagent.aix.rte
```

where `path_name` is the directory path to the SMagent package.

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

```
# lsipp -ah SMagent.aix.rte
```

Starting or stopping SMagent

Use the following command to manually start SMagent:

```
# SMagent start
```

Use the following command to manually stop SMagent:

```
# SMagent stop
```

Configuring 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 6.

Complete the following procedures to configure Storage Manager 8.4 for an AIX system. This can be done from the SMclient running on either an AIX or a non-AIX system.

Adding storage subsystems to SMClient

Perform the following steps to add storage subsystems to SMclient:

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 FAStT storage controllers” on page 17 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 FASStT 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 FASStT storage server. When the page opens, click **Storage Manager v8.4**. Accept the license agreement, then click the **v8.4** link for IBM FASStT 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 FASStT Controller Firmware matrix.

Upgrading FASStT 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. You can grant individual hosts in a defined host group access to storage partitions, independently of the host group. You can make logical drive-to-LUN mappings to the host group or to an individual host in a host group.

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 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.

For more information about storage partitioning, see “Storage Partitioning premium feature” on page 6.

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. If you are configuring additional HBAs to this partition, choose the host port for the next HBA and click **Add**, and then click **Close**.

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

1. Type the following command:

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

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

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

2. Identify the fcs number of the HBA that is connected to the 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.

Mapping LUNs to a FASTt storage partition

Complete the following steps to map LUNs to a storage partition.

Before you begin: Read “Performance planning” on page 3 for information about planning the number of LUNs to map to each partition.

1. Select the Mappings view of the Subsystem Management window.
2. In the topology tree, right-click the host partition on which you want to map LUNs, and select **Define Additional Mappings**. The Define Additional Mapping window opens.
3. In the Define Additional Mapping window, select the following options, and then click **Add**:
 - Host group or host
 - Logical unit number (LUN)(0-255)
 - Logical drive

Repeat these steps for each LUN that you want to map to the partition.

Creating direct and SAN-attached configurations

Storage Manager 8.4 supports IBM FASTt Fibre Channel Storage Servers in direct-attached AIX configurations or in a SAN environment through switches in AIX configurations.

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 servers.
- No external hubs are being used.
- Two-server FAStT700 or FAStT900 configurations require four host-side minihubs, each with exactly one fibre-channel connection from each HBA to a minihub.

Note: FAStT600 storage servers do not have minihubs.

- Single HBA configurations are allowed. If you have a single HBA in a direct-attached configuration, both controllers in the FAStT must be connected to the HBA through a daisy-chain.

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

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

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.

Note: Single HBA configurations are allowed. If you have a single HBA in a SAN-attached configuration, both controllers in the FAStT must be connected to the HBA through a switch, and both controllers must be within the same SAN zone as the HBA.

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

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.

Installing RDAC on AIX hosts

After you install the client software and configure your storage subsystems, use these instructions to install the appropriate Storage Manager 8.4 RDAC device driver.

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

Prerequisites

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

Prerequisites for installing RDAC

See Table 18 on page 119, Table 19 on page 119, or Table 20 on page 120 for a list of required driver version level filesets.

Verifying RDAC software

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

```
# lslpp -ah filename
```

Where *filename* is one of the required file sets listed in Table 18 on page 119, Table 19 on page 119, or Table 20 on page 120.

For example:

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

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 119 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:

```
# lslpp -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
- Viewing Object Data Manager (ODM) attributes in AIX
- Changing ODM attribute settings in AIX
- Verifying the installation and configuration of AIX hosts
- Identifying the controller ID numbers
- Identifying device names and bus numbers
- Identifying logical drives by operating system device names
- Identifying FlashCopy volume types

Viewing and setting attributes of the RDAC driver for AIX

The RDAC driver must be installed on all AIX hosts that will be attached to a FAS*T* storage subsystem.

The RDAC driver creates the following devices that represent the FAS*T* storage subsystem configuration:

- dar** The disk array router represents the entire array, including the current and the deferred paths to all LUNs (hdisk*s* on AIX).
- dac** The disk array controller devices represent a controller within the storage subsystem. There are two dacs in the storage subsystem.
- hdisk** These devices represent individual LUNs on the array.
- 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 FAS*T*.

Note: You might see the utm device listed in command output, whether or not you have an in-band configuration. For example, a utm might be listed when you run the **lsattr** command on a dac.

When you configure these devices, 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 37.

See “Viewing Object Data Manager attributes in AIX” on page 35 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 FAS*T* 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.

Note: Some attributes with False in the Changeable column can be modified using the SMclient.

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 8, Table 9 on page 32, and Table 10 on page 33. Display the default with the **lsattr -dl** command.

Table 8. 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 8. 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 FAStT configuration.
<i>cache_size</i>	Cache size for both controllers, in megabytes; 0 if the sizes do not match.	False	512 or 1024. Set by FAStT.
<i>switch_retries</i>	Number that specifies how many times to retry failed switches, in integers.	True	0 - 255. Should not be changed.

Table 9. Attributes for dac devices

Attribute	Definition	Changeable (T/F)	Possible value
<i>passive_control</i>	Indicator that shows whether this controller was in passive state at the time of configuration (yes) or not (no).	False	Yes or No. State of FAStT configuration.

Table 9. 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 FAST configuration.
<i>controller_SN</i>	Serial number of this controller.	False	Set by FAST.
<i>ctrl_type</i>	Type of array this controller belongs to. A value of value of 1722 indicates FAST600; a value of 1742 indicates FAST700; a value of 1742-900 indicates FAST900.	False	1742, 1722, 1742-900. Set by FAST.
<i>cache_size</i>	Cache size of this controller, in megabytes.	False	512, 1024. Set by FAST.
<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 FAST.
<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 FAST Storage Manager.
<i>node_name</i>	Name of the fibre-channel node.	False	Set by FAST?
<i>location</i>	User-defined location label for this controller; the system does not use this value.	True	Set by FAST Storage Manager.
<i>ww_name</i>	Fibre-channel worldwide name of this controller.	False	Set by FAST.
<i>GLM_type</i>	GLM type used for this controller.	False	High or Low. Set by FAST.

Table 10. Attributes for hdisk devices

Attribute	Definition	Changeable (T/F)	Possible value
<i>pvid</i>	AIX physical volume identifier, or none if not set.	False	Set by AIX.
<i>q_type</i>	Queueing type for this device; must be set to simple .	False	Set by AIX. Must be "simple".

Table 10. Attributes for hdisk devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<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.	False	Yes or No.
<i>size</i>	Size of this volume.	False	Set by FAStT.
<i>raid_level</i>	Number that specifies the RAID level of this device.	False	0, 1, 3, 5. Set by FAStT 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 FAStT Storage Manager.

Table 10. 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. 	False	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.	False	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 (hdisk174).

```
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 FAStT can be modified. This section shows the settings that can be used for best performance and availability.

hdisk attribute settings

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

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 the system is restarted or when **cfgmgr** is run.

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 52.
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 the system is restarted, or when **cfgmgr** is run.

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

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

Use one of the following formulas to determine the maximum queue depth for your system.

FASTT storage systems running appware or snapware 05.4x.xx.xx

On FAStT storage systems that are running appware or snapware 05.4x.xx.xx or later, use the following formula to determine the maximum queue depth:

$$1024 / (\text{number-of-hosts} * \text{LUNs-per-host})$$

For example, a system with four hosts, each with 32 LUNs, would have a maximum queue depth of 8:

$$1024 / (4 * 32) = 8$$

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

```
# chdev -l hdiskX -a queue_depth=8 -P
```

FAStT storage systems running appware or snapware 05.30.xx.xx

On FAStT storage systems that are running appware or snapware 05.30.xx.xx or earlier, use the following formula to determine the maximum queue depth:

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

For example, a system with four hosts, each with 32 LUNs, 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 outlined 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 volumes as shown in the following list:

- Each FAStT600 volume is recognized as a “1722 (600) Disk Array Device”.
- Each FAStT700 volume is recognized as a “1742 (700) Disk Array Device”.
- Each FAStT900 volume is recognized as a “1742-900 Disk Array Device”.

The following example shows the output 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

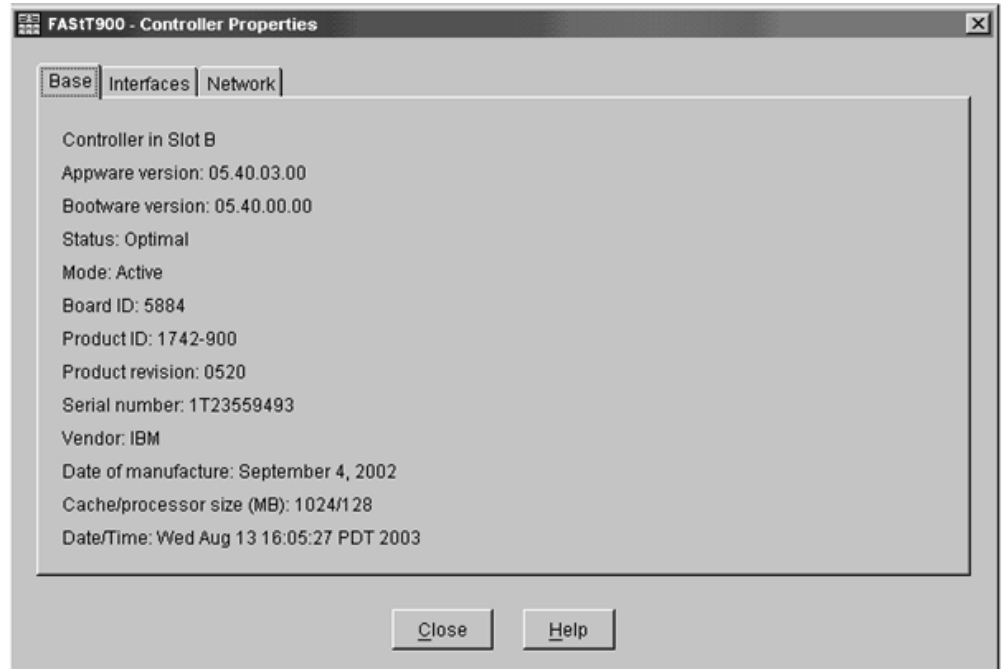
One of the attributes listed by the **lsattr -EI** command is the controller serial number (*controller_SN*) of that dac.

Perform the following step to display the FASTT 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 on page 40 shows the Controller Properties window of the controller in slot A. This controller has a serial number of 1T23559493 and is represented by dac1.

```
# lsattr -EI 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 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
```



SJ001141

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 FAST 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 FAST storage subsystem. `dars` represent the disk array routers; `dacs` represent the disk array controllers, and `hdisks` represent the logical drives, or FAST 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 on page 41.

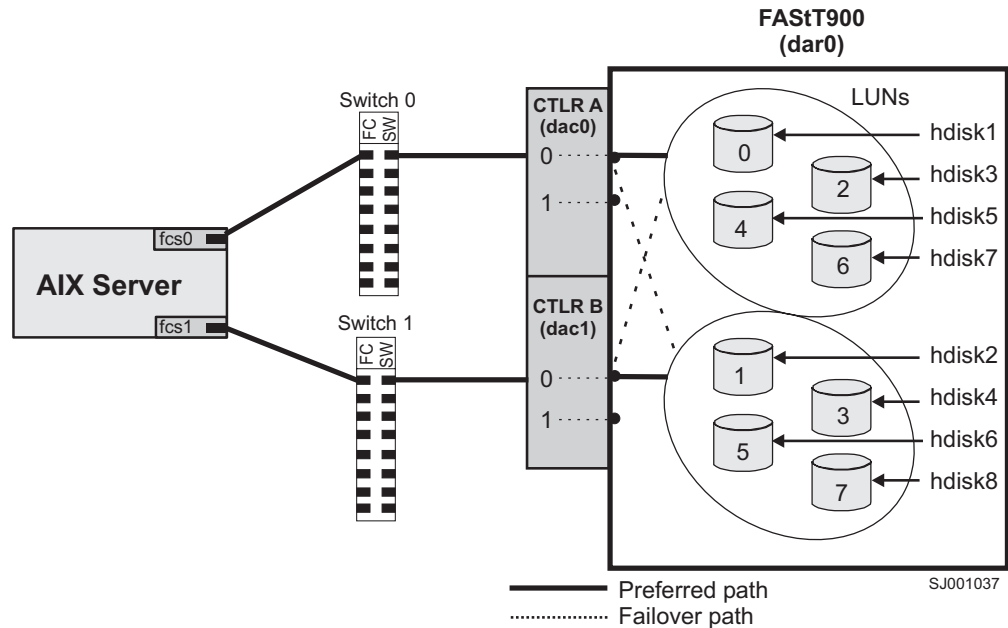


Figure 8. Use of commands applied to validating system configuration

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 dar0 as a machine type 1742, which is a FAST700. 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          1742      (700) Disk Array Router
```

The following example shows the use of the `lsdev` command to view the status of two dacs, which represent the FAST storage subsystem controllers. As in the previous example, a FAST700 is shown 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-adaptor locations. See the service manual for your system type to identify device locations.

```
# lsdev -C |grep dac
dac0      Available 11-08-01      1742      (700) Disk Array Controller
dac1      Available 14-08-01      1742      (700) Disk Array Controller
```

The following example uses the `lsdev` command to show the status and location codes of eight FAST700 hdisks.

Note that the location codes of the odd-numbered hdisks correspond to the location code of dac0 that is listed in the previous example, and that the location codes of the even-numbered hdisks correspond to the location code of dac1. This means that the preferred I/O path for hdisks 1, 3, 5, and 7 is through dac0, and the failover

path is through dac1. Conversely, the preferred path for hdisks 2, 4, 6, and 8 is through dac1, and failover path is through dac0.

```
# lsdev -Cc |grep hdisk
hdisk0    Available 40-60-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1    Available 11-08-01      1742    (700) Disk Array Device
hdisk2    Available 14-08-01      1742    (700) Disk Array Device
hdisk3    Available 11-08-01      1742    (700) Disk Array Device
hdisk4    Available 14-08-01      1742    (700) Disk Array Device
hdisk5    Available 11-08-01      1742    (700) Disk Array Device
hdisk6    Available 14-08-01      1742    (700) Disk Array Device
hdisk7    Available 11-08-01      1742    (700) Disk Array Device
hdisk8    Available 14-08-01      1742    (700) 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
fcsi0     Available 11-08      FC Adapter
fscsi0    Available 11-08-01  FC SCSI I/O Controller Protocol Device
dac0      Available 11-08-01  1742    (700) Disk Array Controller
hdisk1    Available 11-08-01  1742    (700) Disk Array Device
hdisk3    Available 11-08-01  1742    (700) Disk Array Device
hdisk5    Available 11-08-01  1742    (700) Disk Array Device
hdisk7    Available 11-08-01  1742    (700) Disk Array Device
hdisk8    Available 11-08-01  1742    (700) 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 | -I | -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.
- **I** 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 -I 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
(FastT900) #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 FAST900 #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 FAST900 #2'
dac2 ACTIVE dac5 ACTIVE

Disk   DAC   LUN Logical Drive
utm    31
hdisk25 dac5  0 1

---dar3---
User array name = 'D59 FAST900 #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:

```

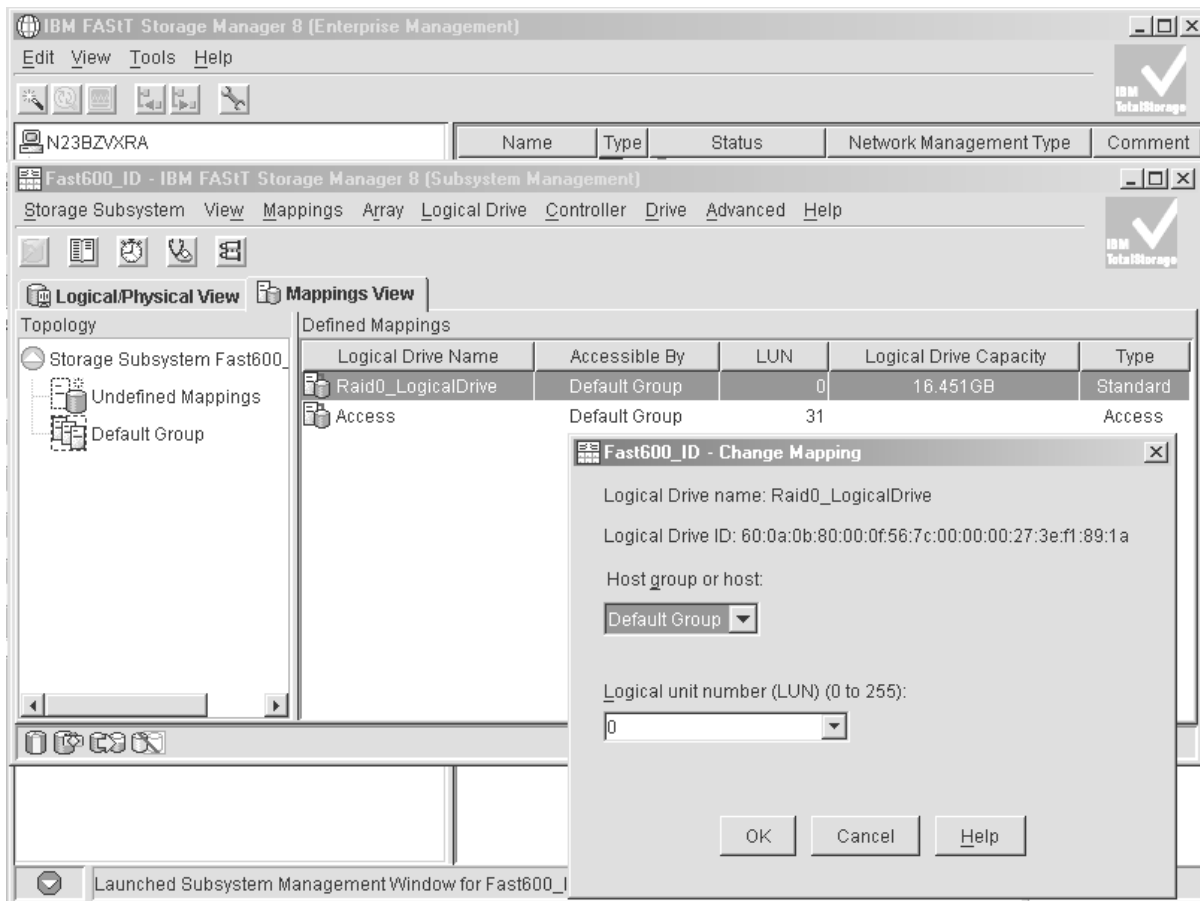
# lsattr -El hdisk4
pvid          none                Physical volume identifier      False
q_type        simple              Queuing Type                   False
queue_depth   30                          Queue Depth                    True
reserve_lock  yes                          RESERVE device on 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 600A0B80000F567C000000273EF1891A 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 600A0B80000F567C000000273EF1891A.

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 *Logical Drive ID* on the FASTt subsystem. You can see the *Logical Drive ID* in the Change Mapping window of the SMclient, which is shown in Figure 9 on page 48.

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**.



SJ001140

Figure 9. Change Mappings window

You can also identify the system disk name by viewing the *Subsystem ID (SSID)* attribute in the Logical Drive Properties window of the SMclient. The value of the *Subsystem ID (SSID)* attribute on the FASTt 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

Perform the following steps to identify FlashCopy volume types:

1. Run the following command to list available volumes:

```
# lsdev -Cc disk
```

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

```
[root@elml7c171] / # 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 1742-900 (900) Disk Array Device
hdisk3 Available 11-08-01 1742-900 (900) Disk Array Device
hdisk4 Available 11-08-01 1742-900 (900) 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 = 'FASTT System 1'
dac0 ACTIVE dac1 ACTIVE
dac0-hdisk2      Raid-0-0B
dac1-hdisk3      Raid-0-1B
dac1-hdisk4      Raid-3-0A-1  Snapshot Volume
```

The FlashCopy volume is identified by the term Snapshot Volume.

For more information about the **lsdev** and **fget_config** commands, see “Identifying device names and bus numbers” on page 40.

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 between the switch and the FASTT storage server. You can change the failover characteristics by setting the fscsi 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 between the switch and the FASTT, 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.

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, the **delayed_fail** setting should be used.

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 fscsi device into a Defined state before you set the attribute.

```
chdev -l fscsi0 -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 **dyntrk** attribute on each HBA that is configured to the FAStT storage server.
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 supported only 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, perform a dynamic capacity expansion (DCE) operation before you perform the DVE operation. A DCE operation increases the capacity of the array by adding physical disks.

Performing a dynamic capacity expansion operation

Complete the following steps to perform a 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. See the man page for the **chvg** command for more information about expanding Logical Volume Manager (LVM) volumes.

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

```
umount mount_point
```

where *mount_point* is the name of the filesystem that is being unmounted.

2. Change the volume group 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.)

3. Mount the file systems.

Notes:

1. You might have to run the **varyoffvg** command, and 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 resize while the volume group is activated in classic or enhanced concurrent mode.
3. Resizing of 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 FAST 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 FAST 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 FAST 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 paths:

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

Note: If a large number of LUNs is configured on the FAStT system, redistributing volumes might take 60 minutes or more to complete, depending on how busy the system is.

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 information for HP-UX operating systems:

- “Hardware and firmware requirements”
- “Installing the client software” on page 56
- “Installing the host software” on page 58
- “Configuring storage subsystems on HP-UX hosts” on page 60
- “Creating direct-attached and SAN-attached configurations” on page 63
- “Enabling multipath I/O with PV-links” on page 64
- “Redistributing volumes in case of failure” on page 69

Attention: For cluster configurations, complete all applicable configuration procedures for each storage subsystem *before* installing the Storage Manager 8.4 software on a second host or cluster server.

Hardware and firmware requirements

For system-level hardware requirements, software requirements, and patch levels, see Appendix B, “HP-UX system requirements,” on page 121.

Attention: Be sure to install the following patches on your HP-UX operating system:

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

Table 11 lists the FAStT Fibre Channel storage servers that are supported with Storage Manager 8.4.

Table 11. FAStT storage servers that are supported on HP-UX systems

Product Name	Model	Minimum appware/snapware (firmware) and NVSRAM versions
IBM TotalStorage FAStT900 Fibre Channel Storage Server	1742-90U	Appware 5.40.06.xx NVSRAM N1742F900R840V02
IBM TotalStorage FAStT700 Fibre Channel Storage Server	1742-1RU	Appware 5.40.06.xx NVSRAM N1742F700R840V03
IBM TotalStorage FAStT600 and FAStT600 Turbo Fibre Channel Storage Servers	1722-60U, 1722-60X	Snapware 5.40.07.xx NVSRAM N1722F600R840V04

Table 12 on page 56 lists the supported versions of other hardware to use with Storage Manager 8.4.

FAStT EXP100

HP-UX operating systems do not support the FAStT EXP100 SATA storage expansion unit.

Table 12. Supported versions of hardware for HP-UX systems

Product Name	Model	Product release and firmware version
IBM TotalStorage FAStT EXP700 Fibre Channel Expansion Unit	1740-1RU	ESM 9325
IBM TotalStorage FAStT EXP500 Fibre Channel Expansion Unit	3560-1RU	ESM 9166
HP FC Adapter	A5158A, A6795A	HP-UX 11.0: B.11.00.10 HP-UX 11.i: B.11.11.09
Brocade switch	2109-S08 2109-S16	2.6.1.b
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 v8.4**. Accept the license agreement, then click the **v8.4** 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 121.
- 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 cdcase -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. Click **Actions** → **Mark for installation**.
5. Click **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. 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 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 121.
- 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. 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 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.

Configuring storage subsystems on HP-UX hosts

After you install the Storage Manager 8.4 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 105 for additional information about storage configuration.

Before you begin: Read the following information:

- All FASTt storage servers have an access volume set to LUN 31.

Note: Do not delete the access LUN or the access volume.

- 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 v8.4**. Accept the license agreement, then click the **v8.4** 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 “Upgrading FASTt firmware and NVSRAM” on page 61. 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 6 for more information.

Complete the following procedures to configure Storage Manager 8.4 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 8.4 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 8.4 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.

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 v8.4**. Accept the license agreement, then click the **v8.4** link for IBM FASSt Storage Manager for HP-UX.

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 firmware:

1. Click **Storage Subsystem** → **Download** → **Firmware**.
2. Go to the directory where the firmware resides.
3. Type or select the full path name 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
```

Upgrading the NVSRAM

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 path name of the NVSRAM file, and then click **Update** → **OK**.

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

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. You can grant individual hosts in a defined host group access to storage partitions, independently of the host group. You can make logical drive-to-LUN mappings to the host group or to an individual host in a host group.

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.

For more information about Storage Partitioning, see “Storage Partitioning premium feature” on page 6.

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 create 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. If the restart does not work, manually enter the WWPNs when prompted.

- 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**.

Mapping LUNs to a FAS*St* storage partition

Before you begin: Read “Performance planning” on page 3 for information about planning the number of LUNs to map to each partition.

Note: HP-UX operating systems do not currently support 256 LUNs per partition.

- HP-UX 11.10 supports up to 32 LUNs per partition.
- HP-UX 11.i supports up to 128 LUNs per partition.

Complete the following steps to map LUNs to a storage partition:

1. Select the Mappings view of the Subsystem Management window.
2. In the topology tree, right-click the host partition on which you want to map LUNs, and select **Define Additional Mappings**. The Define Additional Mappings window opens.
3. In the Define Additional Mappings window, select the following options, and then click **Add**:
 - Host group or host
 - Logical unit number (LUN)
 - The logical drive that you want to map

Repeat these steps for each LUN that you want to map to the partition.

Creating direct-attached and SAN-attached configurations

Storage Manager 8.4 supports IBM FAS*St* Fibre Channel 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 FAS*St* storage servers.
- Two server FAS*St*700 or FAS*St*900 configurations require four host-side minihubs, each with exactly one fibre-channel connection from each HBA to a minihub.

Note: FAS*St*600 storage servers do not have minihubs.

- There are two or four HBAs (HP A6795A or HP A5158A) per FAS*St* storage server. Each pair must be configured to one FAS*St* partition.
- No external hubs are being used.

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

1. Connect the HBAs to each controller or minihub port of the FAS*St* 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 FAS*St* 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.

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.

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 65. 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 66.

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 66, you notice that they differ in only five digits, f56d0 and f5d6c.

The devices in Figure 11 on page 66 are viewed through the controllers c166 and c172. To determine the preferred path of a specific logical drive seen by the operating system perform 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 13).

Table 13. 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 step 3a on page 67 through step 3c on page 67 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, perform the following steps to create volumes and volume groups.

Note: The steps in this procedure refer to LUN 0 in Table 13 on page 67.

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 step 1 on page 68 through step 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:

```
#vgdisplay -v vgname
```

where *vgname* is the volume group name.

Redistributing volumes in case of failure

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

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

1. Repair or replace any faulty components. For more information, see the *IBM TotalStorage FAStT Installation and Support Guide* for the appropriate FAStT 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 for Solaris operating systems:

- “Hardware and firmware requirements”
- “Installing the client software” on page 73
- “Installing the host software” on page 75
- “Configuring storage subsystems on Solaris hosts” on page 77
- “Creating direct-attached and SAN-attached configurations” on page 79
- “Configuring the host bus adapter cards” on page 80
- “Setting up a Solaris host group” on page 88
- “Mapping LUNs to a FAStT storage partition” on page 89
- “Installing RDAC on Solaris hosts” on page 89
- “Installing DMP on Solaris hosts” on page 92
- “Verifying external storage” on page 102
- “Redistributing volumes in case of failure” on page 103

Attention: For cluster configurations, complete all applicable configuration procedures for each storage subsystem *before* installing the Storage Manager 8.4 software on a second host or cluster server.

Hardware and firmware requirements

Table 14 lists the FAStT Fibre Channel storage servers that are supported with Storage Manager 8.4.

Table 14. FAStT storage servers that are supported on Solaris systems

Product name	Model	Minimum appware/snapware (firmware) and NVSRAM versions
IBM TotalStorage FAStT900 Fibre Channel Storage Server	1742-90U	Appware 5.40.06.xx NVSRAM N1742F900R840V02
IBM TotalStorage FAStT700 Fibre Channel Storage Server	1742-1RU	Appware 5.40.06.xx NVSRAM N1742F700R840V03
IBM TotalStorage FAStT600 and FAStT600 Turbo Fibre Channel Storage Server	1722-60U, 1722-60X	Snapware 5.40.07.xx NVSRAM N1722F600R840V04

Table 15 on page 72 lists the supported versions of other hardware that are supported with Storage Manager 8.4.

FAStT EXP100

Solaris operating systems do not support the FAStT EXP100 SATA storage expansion unit.

Table 15. Supported versions of hardware for Solaris systems

Product Name	Model	Product release and firmware version
IBM TotalStorage FAStT EXP700 Fibre Channel Expansion Unit	1740-1RU	ESM 9325
IBM TotalStorage FAStT EXP500 Fibre Channel 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 4.1.3
JNI FC Adapter	FCE-1473 S-bus 64 bit	JNIC146x.pkg, version 5.2.3
JNI FC Adapter	FCE2-1473 S-bus 64 bit	JNIC146x.pkg, version 5.2.3
JNI FC Adapter	FCE-6410-N 33 Mhz	JNIC.pkg, version 4.1.3
JNI FC Adapter	FCE-6460-PCI 64 bit	JNIC146x.pkg, version 5.2.3
JNI FC Adapter	FCE2-1063 S-bus 64 bit	JNIC.pkg, version 4.1.3
JNI FC Adapter	FCE2-6412 66 Mhz 64 bit	JNIC.pkg, version 4.1.3
JNI FC Adapter	FCI-1063	fca-pci.2.5.18.pkg
JNI FC Adapter	FCC-6562	JNIC146x.pkg, version 5.2.3
JNI FC Adapter	FCC2-6562	JNIC146x.pkg, version 5.2.3
JNI FC Adapter	FCX-6562	JNIC146x.pkg, version 5.2.3
JNI FC Adapter	FCC\X2-6562	JNIC146x.pkg, version 5.2.3
QLogic Adapter	4.08	2.00.01
Brocade switch	2109-S08, 2109-S16	2.6.1.b
Brocade switch	2109-F16	3.1.1.b
Brocade switch	2109-F8	3.0.2.K
Brocade switch	2109-M12, 2109-F32	4.1.1.a
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 v8.4**. Accept the license agreement, then click the **v8.4** link for IBM FAStT Storage Manager for Solaris.

For system-level hardware and software requirements and patch levels, see Appendix C, “Solaris system requirements,” on page 123.

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 7 or 8 and meets the minimum hardware and software requirements described in Appendix C, “Solaris system requirements,” on page 123.
- 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 reboot is required during the client installation process.

Before you begin: Visit the Web sites that are listed in “Hardware and firmware requirements” on page 71 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/SM84/SUN_SParc/SM8runtime-Sparc-filename.pkg
```

where *filename* is the name of the package that you want to install. 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:

```
# pkginfo -l 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 7 and 8.

Adjust these instructions as required for your specific installation location. No reboot 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/SM84/SUN_Sparc/SM8client-Sparc-filename.pkg
```

where *filename* is the name of the package that you want to install. 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:

```
# pkginfo -l 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 74 through step 6 of the installation procedure. Use the corresponding installation profile for each storage management station as a guide.

Installing the host software

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)
- Failover driver (RDAC or DMP)

This section contains the procedure for installing SMagent and SMutil. For information about how to install RDAC and DMP, see “Installing RDAC on Solaris hosts” on page 89 and “Installing DMP on Solaris hosts” on page 92.

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 7 or 8, and meets the minimum hardware and software requirements described in Appendix C, “Solaris system requirements,” on page 123.
- The SMagent, SMutil, and RDAC 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 steps to install SMagent:

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

```
pkgadd -d /cdrom/SM84/SUN_SParc/SM8agent-Sparc-filename.pkg
```

where *filename* is the name of the package that you want to install.
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:

```
# pkginfo -l 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 ./pathname/SMutil-Sparc-filename.pkg
```

where *filename* is the name of the package that you want to install.
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:

```
# pkginfo -l SMutil
```

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.

Configuring storage subsystems on Solaris hosts

To configure Storage Manager 8.4 for a Solaris system, complete the following procedures in the order that they are described in this section:

- Add storage subsystems to SMclient
- Update NVSRAM and FAStT firmware (if necessary)

Before you begin: Read the following information:

- All FAStT storage servers have an access volume set to LUN 31.

Note: Do not delete the access LUN or the access volume.

- If you have purchased a Storage Partitioning premium feature, ensure that the feature is enabled. See “Enabling your premium features” on page 6 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 v8.4**. Accept the license agreement, then click the **v8.4** 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 “Upgrading 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**.

Upgrading 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 v8.4**. Accept the license agreement, then click the **v8.4** 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.

Upgrading the NVSRAM

Complete the following steps to upgrade the NVSRAM:

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 path name of the NVSRAM file, and then click **Update** → **OK**.

Upgrading FASTT firmware

Complete the following steps to upgrade the firmware:

1. Click **Storage Subsystem** → **Download** → **Firmware**.
2. Type or select the full path name 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.

Creating direct-attached and SAN-attached configurations

Storage Manager 8.4 supports IBM FAStT Fibre Channel 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 FAStT storage servers.

Requirements:

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

Note: FAStT600 storage servers do not have minihubs.

- 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.

Configuring the host bus adapter cards

Before you can install the failover driver, you must install the HBA driver package and modify the configuration file settings using the following procedures:

JNI

- “Installing the JNI adapter driver package”
- “Modifying the JNI HBA settings” on page 81 (SAN-attached only)

QLogic

- “Installing the QLogic adapter driver package” on page 86
- “Modifying the QLogic HBA settings” on page 87

Note: The QLogic HBA is supported only on Solaris 8 operating systems.

Installing the JNI adapter driver package

Before you begin:

- Ensure that the JNI HBAs have been installed and attached directly to the controllers or to a SAN fabric switch.
- If attached to SAN fabric switch, ensure that the zones for the FAStT storage subsystem have been created and enabled. For information about enabling zones, see the documentation that is provided by the switch manufacturer.

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 16.

Table 16. JNI model numbers, adapter drivers, and configuration files

JNI HBA model numbers	JNI adapter driver package	Configuration file
FC64-1063-N	fcaw2.5.18.pkg	/kernel/drv/fcaw.conf
FCI-1063-N	fca-pci.2.5.18.pkg	/kernel/drv/fca-pci.conf
FCE2-1063, FCE2-1063, FCE-6410, FCE2-6412	JNIC.pkg, version 4.1.1.1	/kernel/drv/jnic.conf
FCE-1473, 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 16, 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 89 or “Installing DMP on Solaris hosts” on page 92. 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 83 for the correct modification procedures for a Brocade 2 Gb switch.

Before you begin: Table 16 on page 80 lists the configuration file for each JNI adapter driver package. Also, some of the steps in these procedures refer to Appendix E, “JNI and QLogic host bus adapter settings,” on page 127.

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 and QLogic host bus adapter settings,” on page 127.
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.

```
# /SunOS
```

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 81 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

```

- 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.

```

- 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 and QLogic host bus adapter settings,” on page 127 and modify settings accordingly.

- Type the following command:

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

- 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 84, 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” on page 89 or “Installing DMP on Solaris hosts” on page 92 for more information.

Installing the QLogic adapter driver package

This section describes how to install the QLogic adapter driver package.

Note: The QLogic HBA is supported only on Solaris 8 operating systems.

Before you begin:

- Ensure that the QLogic HBAs have been installed, and are attached either directly to the controllers or to the SAN fabric switch.
- If the HBAs are attached to a SAN fabric switch, ensure that the zones for the FASTT storage subsystem are created and enabled. For information about enabling zones, see the documentation that is provided by the switch manufacturer.

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

Complete the following steps to install the QLogic adapter driver package.

1. Download the most current adapter driver package from the following Web site to a local directory on your host:

www.QLogic.com

2. Uncompress the QLogic adapter driver package, by typing the following command:

```
#uncompress qllogic_adapter_driver.Z
```

where *qllogic_adapter_driver* is the name of the compressed adapter driver (for example, *qla2300_pkg_v408.Z*).

3. Install the QLogic adapter driver package, by typing the following command:

```
#pkgadd -d adapter_driver_pkg
```

where *adapter_driver_pkg* is the name of the package that you want to install (for example, *qla2300_pkg_408*). A package installation menu displays.

4. From the package installation menu, select package 3:

```
3. QLA2300-3 QLogic QLA23000 driver
   (Sparc) Solaris 8-9, Rev=4.08
```

Follow the instructions in the package installation utility to install the QLogic adapter driver package.

Note: You do not need to install any of the QLSDMLIB packages that display in the package installation menu.

5. Verify that the QLogic adapter package is installed, by typing the following command:

```
#pkginfo -l pkg_version
```

where *pkg_version* is the name of the adapter driver package that you installed (for example, *QLA2300-3*).

6. Reboot the Solaris host, by typing the following command:

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

Modifying the QLogic HBA settings

Complete the following steps to modify the QLogic HBA settings.

1. Open the QLogic configuration file, by typing the following commands:

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

where *QLogic_conf_file.conf* is the name of the QLogic configuration file (for example, *qla2300.conf*).

2. In the Vi Editor, uncomment and modify the loop attributes using the information in Appendix E, “JNI and QLogic host bus adapter settings,” on page 127.
3. Save the changes to the QLogic configuration file, by typing the following command:

```
#:wq
```

4. Reboot the Solaris host, by typing the following command:

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

When you are finished, install the failover driver on the host. See “Installing RDAC on Solaris hosts” on page 89 or “Installing DMP on Solaris hosts” on page 92 for more information.

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. You can grant individual hosts in a defined host group access to storage partitions, independently of the host group. You can make logical drive-to-LUN mappings to the host group or to an individual host in a host group.

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 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.

For more information about storage partitioning, see “Storage Partitioning premium feature” on page 6.

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. Click **File** → **Exit**.

Mapping LUNs to a FAS*T* storage partition

Before you begin: Read “Performance planning” on page 3 for information about planning the number of LUNs to map to each partition.

Complete the following steps to map LUNs to a storage partition:

1. Select the Mappings view of the Subsystem Management window.
2. In the topology tree, right-click the host partition on which you want to map LUNs, and select **Define Additional Mappings**. The Define Additional Mappings window opens.
3. In the Define Additional Mappings window, select the following options, and then click **Add**:
 - Host group or host
 - Logical unit number (LUN)
 - The logical drive that you want to map

Repeat these steps for each LUN that you want to map to the partition.

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 80.

For information about how to set the bindings, see “Modifying the JNI HBA settings” on page 81.

- 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 v8.4**. Accept the license agreement, then click the **v8.4** 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: In this procedure, the installation CD is mounted at /cdrom/SM84. Adjust the procedure as required for your specific installation location.

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

```
# pkgadd -d /cdrom/SM84/SUN_Sparc/SM8rdac-Sparc-filename.pkg
```

where *filename* is the name of the package on the Solaris installation CD that you want to install.

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 81.

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 steps, in which you remove RDAC and reboot the host.

Note: You do not need to remove RDAC or reboot the host when you modify or add targets or WWPNs for the controllers. New HBAs can bind to controllers without removing RDAC or rebooting the host. To modify or add targets or WWPNs for controllers, do not complete the following steps; 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 by typing the following command:

```
#pkgrm RDAC_driver_pkg_name
```

where *RDAC_driver_pkg_name* is the name of the RDAC driver package that you want to uninstall.

2. Verify RDAC drive package removal by typing the following command:

```
#pkginfo RDAC_driver_pkg_name
```

where *RDAC_driver_pkg_name* is the name of the RDAC driver package that you uninstalled.

3. Reboot the Solaris host by typing the following command:

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

4. Modify the JNI configuration file by typing the following command:

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

where *JNI_configuration_file* is the name of the JNI configuration file that you want to modify. When you have finished making changes, run the following command to save the changes:

```
#:wq
```

5. Install the RDAC driver package by typing the following command:

```
#pkgadd -d RDAC_driver_pkg_name
```

where *RDAC_driver_pkg_name* is the name of the RDAC driver package that you want to install.

6. Verify package installation by typing the following command:

```
#pkginfo RDAC_driver_pkg_name
```

where *RDAC_driver_pkg_name* is the name of the RDAC driver package that you installed.

7. Reboot the Solaris host by typing the following command:

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

Installing DMP on Solaris hosts

This section describes how to install VERITAS Dynamic Multipathing (DMP), which is a failover driver for Solaris hosts. While RDAC supports only 32 LUNs, DMP supports up to 256 LUNs.

Important: For more information about using VERITAS DMP, see the VERITAS Volume Manager and VERITAS File System v3.5 documentation for Solaris, which you can download from the following Web site:

www.veritas.com

System requirements

Ensure that your system meets the following requirements for installing DMP:

- Solaris operating system
- VERITAS Volume Manager 3.5
- SMibmasl.pkg (enables Solaris to recognize the FAStT machine type)

DMP installation overview

Ensure that your system meets the following prerequisites for installing DMP:

- The HBAs are installed on the Solaris host.
- The JNI HBA package is installed.
- The parameter settings in the JNI configuration file (for example, Jnic146x.conf) are modified.
- In a SAN environment, bindings are configured.
- The zones are created and enabled for the Solaris partition.
- Storage is mapped to the Solaris partition.

Perform the following procedures, in the order listed, to complete the DMP installation:

1. "Preparing for VERITAS DMP installation"
2. "Installing VERITAS Volume Manager packages for DMP" on page 94
3. "Installing the SMibmasl package" on page 97
4. "Configuring VERITAS Volume Manager" on page 98
5. "Starting the VERITAS Enterprise Administrator" on page 101

Preparing for VERITAS DMP installation

Complete the following steps to prepare the host for installing VERITAS DMP:

1. Choose the Solaris host on which you want to install DMP.
2. Manually define the targets and LUNs in the `/kernel/drv/sd.conf` file, by completing the following steps.

By default, the `/kernel/drv/sd.conf` file defines targets 0, 1, 2, and 3. LUN0 also is defined for targets 0, 1, 2, and 3.

Notes:

- Each target represents a controller to a subsystem, and each LUN represents a logical volume.
- If you are adding additional target or LUN definitions to the `/kernel/drv/sd.conf` file for an existing DMP configuration, be sure to reboot the Solaris host.

- a. Open the `/kernel/drv/sd.conf` file with the Vi Editor, by typing the following command:

```
# vi /kernel/drv/sd.conf
```

The file will look similar to the following example:

```
#
# Copyright (c) 1992, Sun Microsystems, Inc.
#
#ident "@(#)sd.conf 1.9 98/01/11 SMI"

name="sd" class="scsi" class_prop="atapi"
target=0 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=1 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=2 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=3 lun=0;
```

- b. Add additional target and LUN definitions, using the Vi Editor. In the following example, it is assumed that the Solaris host is attached to one FASTT subsystem with three LUNs mapped to the FASTT storage partition. In addition, the access LUN must be mapped to the partition.

```
#
# Copyright (c) 1992, Sun Microsystems, Inc.
#
#ident "@(#)sd.conf 1.9 98/01/11 SMI"

name="sd" class="scsi" class_prop="atapi"
target=0 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=1 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=2 lun=0;

name="sd" class="scsi" class_prop="atapi"
target=3 lun=0;

name="sd" class="scsi" target=0 lun=1;
name="sd" class="scsi" target=0 lun=2;
name="sd" class="scsi" target=0 lun=3;
name="sd" class="scsi" target=0 lun=31;
name="sd" class="scsi" target=1 lun=1;
name="sd" class="scsi" target=1 lun=2;
name="sd" class="scsi" target=1 lun=3;
name="sd" class="scsi" target=1 lun=31;
```

- c. Save the new entries in the `/kernel/drv/sd.conf` file, by typing the following command:

```
#:wq
```

3. Verify that RDAC is not installed on the host, by typing the following command:

```
# pkg info -l RDAC
```

4. If RDAC is installed, remove it by typing the following command:

```
# pkgrm RDAC
```

5. Verify that a host partition has been created. If not, create one by following the procedure that is described in “Setting up a Solaris host group” on page 88.

Note: Set the host port type to Solaris DMP. Failure to do so results in an inability to map for more than the RDAC limit of 32 LUNs, and causes other undesired results.

6. Ensure that all of the paths are optimal, and are in a preferred path state from the SMclient.
7. Reboot the host and ensure that the devices show up, by typing the following command:

```
# SMdevices
```

8. Verify that you have a 1 GB logical drive available for the rootdg.
9. Label the logical drive that you are using for the rootdg, by typing the following command:

```
# format
```

10. Reboot the Solaris host, by typing the following command:

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

Installing VERITAS Volume Manager packages for DMP

The installation steps that are described in this section are based on a CD installation of VERITAS.

Before you begin: Ensure that you have the VERITAS Volume Manager license key, and any other VERITAS licenses that you want to install, available during the installation process.

Complete the following steps to install VERITAS Volume Manager packages for DMP:

1. Insert the VERITAS CD into the CD-ROM drive.
2. Navigate to the mounted CD directory. If you are unsure about the mount point path to the CD-ROM, type the following command:

```
#df
```

The result is similar to the following example:

```

/          (/dev/dsk/c0t0d0s0):22222222 blocks 173544 files
/proc     (/proc                ):          0 blocks 29903 files
/dev/fd   (fd                  ):          0 blocks    0 files
/etc//mnttab (mnt                ):          0 blocks    0 files
/var/run  (swap                 ): 6764416 blocks 210436 files
/tmp      (swap                 ):6764416 blocks 210436 files
/cdrom/storage_solutions_solaris_3.5cd1
(/vol/dev/dsk/c1t6d0/storage_solutions_solaris_3.5cd1):
22222222 blocks 500 files

```

Note: In this example, you would type the following command to go to the mounted CD directory:

```
# cd /cdrom/storage_solutions_solaris_3.5cd1
```

3. Run the VERITAS installer script, by typing the following commands:

```
# ls
```

Note: The output of the **# ls** command is similar to the following:

```

cluster_server      menu
database_ac_for_oracle9i sanpoint_foundation_suite
file_system         traffic_director
installer           volume_manager
installer.pdf       win32

```

```
# ./installer
```

4. A new screen displays. Type y (yes) to install VRTSvlic, as shown in the following example:

```

VERITAS Storage Solutions
Looking for package VRTSvlic
Currently installed 0
Minimum Version:3.00.000

For release Train installation to continue, VRTSvlic
must installed or upgraded.
Do you want to install it now? [y,n]: y

```

5. The VERITAS Storage Solutions screen displays. Install the license for VERITAS Volume Manager 3.5 by selecting option 1 (Add License Key), as shown in the following example:

```

VERITAS Storage Solutions
VERITAS Product      Currently Installed  Licensed
=====
Volume Manager      Not Installed      No
File System         Not Installed      No
Foundation Suite    Not Installed      No
Volume Replicator   Not Installed      No
Flash Snap          Not Installed      No
Database Edition for Ora Not Installed      No

1= Add License Key 2=Installation Menu 3=Refresh h=help
p= Product Description q=quit

Enter [1,2,3,h,p,q]: 1

```

6. Type the VERITAS license key, as shown in the following example:

```

VERITAS Storage Solutions
Please enter your VERITAS license key: XXXX-XXXX-XXXX-XXXX-XXXX

```

7. After the license is successfully installed, a product summary screen displays. Install VERITAS Volume Manager by selecting option 2 (Installation Menu), as shown in the following example.

Note: Allow VERITAS to install conflicting files. At minimum, you must install the VRTSvmpro file. If you want to use the GUI version of VERITAS Volume Manager, you must also install additional packages. See the VERITAS Volume Manager 3.5 installation documentation if you want to install the VERITAS Volume Manager GUI.

```

VERITAS Storage Solutions
VERITAS Product      Currently Installed  Licensed
=====
Volume Manager      Not Installed      Yes
File System         Not Installed      Yes
Foundation Suite    Not Installed      Yes
Volume Replicator   Not Installed      No
Flash Snap          Not Installed      No
Database Edition for Ora Not Installed      No

1= Add License Key 2=Installation Menu 3=Refresh h=help
p= Product Description q=quit

Enter [1,2,3,h,p,q]: 2

```

8. A list of available products displays. Install VERITAS Volume Manager by selecting option 1 (Volume Manager) from the products list, then type y (yes), as shown in the following example:

```

VERITAS Storage Solutions
Available products:

1) Volume Manager
2) File System
3) Foundation Suite (Quick Start, HA & Flash Snap)
4) Volume Replicator
5) Database Edition for Oracle
6) Cluster Server
7) SanPoint Foundation Suite
8) Cluster Server QuickStart
9) Database Edition/Advance Cluster for Oracle 9i
10) Cluster Server Traffic Director

q) Return to the main menu
h) Installation help

Enter the number of the product to install or [q,h]: 1

VERITAS Storage Solutions

You are about to begin installation of Volume Manager

Do you want to continue? [y,n]: y

```

9. A list of available products displays again. Install VERITAS file systems by selecting option 2 (File System), then typing y (yes).

Note: Allow VERITAS to install conflicting files, and when prompted to install setuid/setgid files and the VRTSvxf file, answer y (yes) to both.

10. Exit the VERITAS installer by selecting option q (Return to the main menu) on the products screen.

11. A product summary screen displays. Exit the screen by selecting option q (quit), as shown in the following example:

```

VERITAS Storage Solutions
VERITAS Product      Currently Installed  Licensed
=====
Volume Manager      3.5, REV=06.21.2002.23.14 Yes
File System         3.5, REV=GA06      Yes
Foundation Suite    3.5                Yes
Volume Replicator   3.5, REV=06.21.2002.23.14 Yes
Flash Snap          3.5, REV=06.21.2002.23.14 Yes
Database Edition for Ora Not Installed  No

1= Add License Key 2=Installation Menu 3=Refresh h=help
p= Product Description  q=quit

Enter [1,2,3,h,p,q]: q

```

Installing the SMibmasl package

Complete the following steps to install the SMibmasl package:

1. Install the SMibmasl package, by typing the following command:

Note: You can select either the default (all), or select option 1.

```
#pkgadd -d SMibmasl_pkg
```

2. Reboot the Solaris host, by typing the following command:

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

Configuring VERITAS Volume Manager

Complete the following steps to configure the VERITAS Volume Manager:

1. Initialize one disk for rootdg, by typing the following command:

```
# vxinstall
```

Note: You have already entered a license key for VERITAS Volume Manager and VERITAS File System. You do not need to install additional VERITAS software at this time. Unless you want to install additional VERITAS software at this time, type n (no) in response to both of the following questions:

```
Some licenses are already installed.
Do you wish to review them?
[y,n,q,?] (default: y) n

Do you wish to enter another
license key [y,n,q,?]
(default: n) n
```

2. The Volume Manager Installation screen displays. Choose whether to set up enclosure-based names for the FAStT-attached disks by typing y (yes) or n (no), as shown in the following example:

```
Volume Manager Installation
Menu: Volume Manager/Install

VxVM will use the following format to name disks on the host:

enclosurename_diskno

In the above format, enclosurename is the logical name
of the enclosure to which the disk belongs.
VxVM assigns default enclosure names which
can be changed accordingly to the user requirements.

Some examples would be:

FAStT_2 - second disk detected in enclosure 'hitachio'
Jbod_1 - first disk detected in the jbod category

For non-public loop disks (Ex. FC disks directly
connected to the host or through Hubs),
you could choose to use the c#t#d#s# naming format.

Do you want to use enclosure based names for all disks?
[y,n,q,?] (default: n) n
```

3. The Volume Manager Installation screen displays. Press the Return key to continue the installation process, as shown in the following example:

```
Volume Manager Installation
Menu: Volume Manager/Install

    The Volume Manager has detected the following
categories of storage connected to your system:

    Enclosure: Disk   IBM_FASt0

Hit RETURN to continue
```

4. The Volume Manager Installation Options screen displays. Select option 2 (Custom Installation), as shown in the following example.

```
Volume Manager Installation Options
Menu: VolumeManager/Install

1. Quick Installation
2. Custom Installation
3. Prevent multipathing/Suppress devices from VxVM's view

? Display help about menu
?? Display help about the menuing system
q Exit from menus

Select an operation to perform: 2
```

5. The Volume Manager Custom Installation screen displays. Select n (do not encapsulate the boot disk), as shown in the following example:

```
Volume Manager Custom Installation
Menu: VolumeManager/Install/Custom

The CxTxDx disk is your boot disk.
You can add it as a new disk.
If you encapsulate it, you will
make your root file system and
system areas on the Boot Disk
into volumes. This is required if
you wish to mirror your root file
system or system swaparea.

Encapsulate Boot Disk [y,n,q,?] (default: n) n
```

6. The Volume Manager Custom Installation screen displays. Press the Return key to continue, as shown in the following example:

```
Volume Manager Custom Installation
Menu: VolumeManager/Install/Custom/Disk

Disk array serial number: DISKS

Generating list of disks in disk array Disk ...

<excluding root disk c0t0d0>

    The volume Manager has detected the following disks in
Enclosure Disk:

    c0t1d0

Hit RETURN to continue.
```

7. The Installation Options for Controller Disk screen displays. Select option 3 (Install one disk at a time), as shown in the following example:

```
Installation Options for Controller Disk
Menu: Volume Manager/Install/Custom/Disk

 1 Install all disks as pre-existing disks. (encapsulate)
 2 Install all disks as news disks. (discard data on disks!)
 3 Install one disk at a time.
 4 Leave these disks alone.

? Display help menu
?? Display help about the menuing system
q Exit from menus

Select an operation to perform: 3
```

8. The Installation Options for Disk XXXXXX screen displays (similar to the screen that is shown in step 7). Select option 3 (Install one disk at a time).
9. The Volume Manager Custom Installation screen displays. Press the Return key to continue, as shown in the following example:

```
Volume Manager Custom Installation
Menu: VolumeManager/Install/Custom/disk/IBM_FAST0

Disk array serial number : 600A0B800000F14400000

Generating list of disk in disk array IBM_FAST0 ....

    The Volume Manager has detected the following
disk in enclosure IBM_FAST0:

    C8t4d0  c8t1d1  c8t4d2  c8t4d3

Hit RETURN to continue.
```

10. The Installation Options for Controller YYYYYY displays (similar to the screen that is shown in step 7). Select option 3 (Install one disk at a time).
11. The Installation Options for Disk ZZZZZZ displays. For the rootdg, select option 3 (Install as a new disk).

In the following example, c8t4d0 is the logical volume that is being used as the rootdg for VERITAS DMP:

```
Installation options for disk c8t4d0
Menu: VolumeManager /Install/Custom/Disk/IBM_FAST0

 1 Install as pre-existing disk. (encapsulate)
 2 Install as a new disk. (discard data on disks!)
 3 Leave this disk alone.

? Display help menu
?? Display help about the menuing system
q Exit from menus

Select an operation to perform: 2

Are you sure (destroys data on c8t4d0) [y,n,q,?]
  (default: n) Y

Enter disk name for c8t4d0 [<name>,q,?] (default: disk 01)
```


12. The Installation Options for Disk `XYXYXY` displays. Select option 3 (Leave this disk alone) for all non-rootdg logical volumes, as shown in the following example:

```
Installation options for disk c8t4d1
Menu: VolumeManager /Install/Custom/Disk/IBM_FAStT0

 1 Install as pre-existing disk. (encapsulate)
 2 Install as a new disk. (discard data on disks!)
 3 Leave this disk alone.

 ? Display help menu
 ?? Display help about the menuing system
 q Exit from menus

Select an operation to perform: 3
```

13. Be sure to repeat step 12 for each of the remaining disks on the FAStT storage array.
14. A summary screen displays, similar to the following example. Press Enter.

```
Volume Manager Custom Installation
Menu: VolumeManager /Install/Custom/Disk/IBM_FAStT0

The following is a summary of your choices.

C8t4d0  New disk

Is this correct [y,n,?] (default: y) ENTER
```

15. Status information displays, similar to the following example. When the display has completed, select `y` (yes) to reboot the system.

```
The Volume Manager is now reconfiguring (partition phase) ...
Volume Manager: Partitioning c8t4d0 as a new disk
The Volume Manager is now reconfiguring (initialization phase) ...
Volume Manager: Adding disk01 (c8t4d0) as new disk.
The volume Daemon has been enabled for transactions.
The system now must be shut down and rebooted in order to
continue the reconfiguration.
Shutdown and reboot now [y,n,q,?] (default:n) Y
```

Starting the VERITAS Enterprise Administrator

After you have installed the VERITAS packages, and configured the rootdg, you must stop and restart the VERITAS Enterprise Administrator (VEA).

Note: VEA starts automatically after a reboot.

Complete the following steps to start VEA:

1. Stop the VEA service, by typing the following command at the command prompt:

```
# /opt/VRTSob/bin/vxsvc -k
```

- Restart the VEA service, by typing the following command:

```
# /opt/VRTSob/bin/vxsvc
```

- Verify the state of the VEA service, by typing the following command:

```
# /opt/VRTSob/bin/vxsvc -m
```

- Start the VEA GUI, by typing the following command:

```
# /opt/VRTSob/bin/vea &
```

- Login to the VEA GUI, and rescan.

When you are done, you have completed installing VERITAS DMP for the Solaris host.

See the VERITAS documentation for information about how to complete the following tasks:

- Set up disk groups
- Create volumes
- Create file systems
- Mount file systems

Verifying external storage

After you install the failover driver, complete the following steps 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 76.

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.

- 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
```

- Run the following command to cross-verify between SMutil 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 FAStT 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 FAStT Installation and Support Guide* for the appropriate FAStT storage server.
2. Redistribute volumes to their preferred paths by clicking **Subsystem Management** → **Storage Subsystem** → **Redistribute Logical Drive**.

Chapter 5. Completing the software configuration

After you finish installing Storage Manager 8.4, you must complete your configuration. This chapter describes how to complete the Storage Manager 8.4 configuration by performing the following tasks:

- “Configuring storage subsystems”
- “Configuring alert notifications” on page 107
- “Renaming storage subsystems” on page 107
- “Creating a storage subsystem profile” on page 107
- “Performing other storage subsystem management tasks” on page 108

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

Overview of heterogeneous host configurations

The heterogeneous hosts feature enables hosts that are running different operating systems to access a single storage subsystem. Storage Manager 8.4 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 run on 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 you define a host type in the Define New Host Port window, the heterogeneous hosts feature enables the controllers in the storage subsystem to tailor their behavior (such as LUN reporting and error conditions) to the needs of the operating system or variant of the host that is sending the information.

Important: When you define host types in a heterogeneous environment, you must set each host type to its appropriate operating system so that the firmware on each controller can respond correctly for the appropriate operating system.

Before you begin setting up the configuration for you heterogeneous host, see the *IBM FAS*™*T Storage Manager Concepts Guide*.

Configuring storage subsystems

Use the following procedures to start the SMclient 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 SMclient to discover and configure each attached storage subsystem.

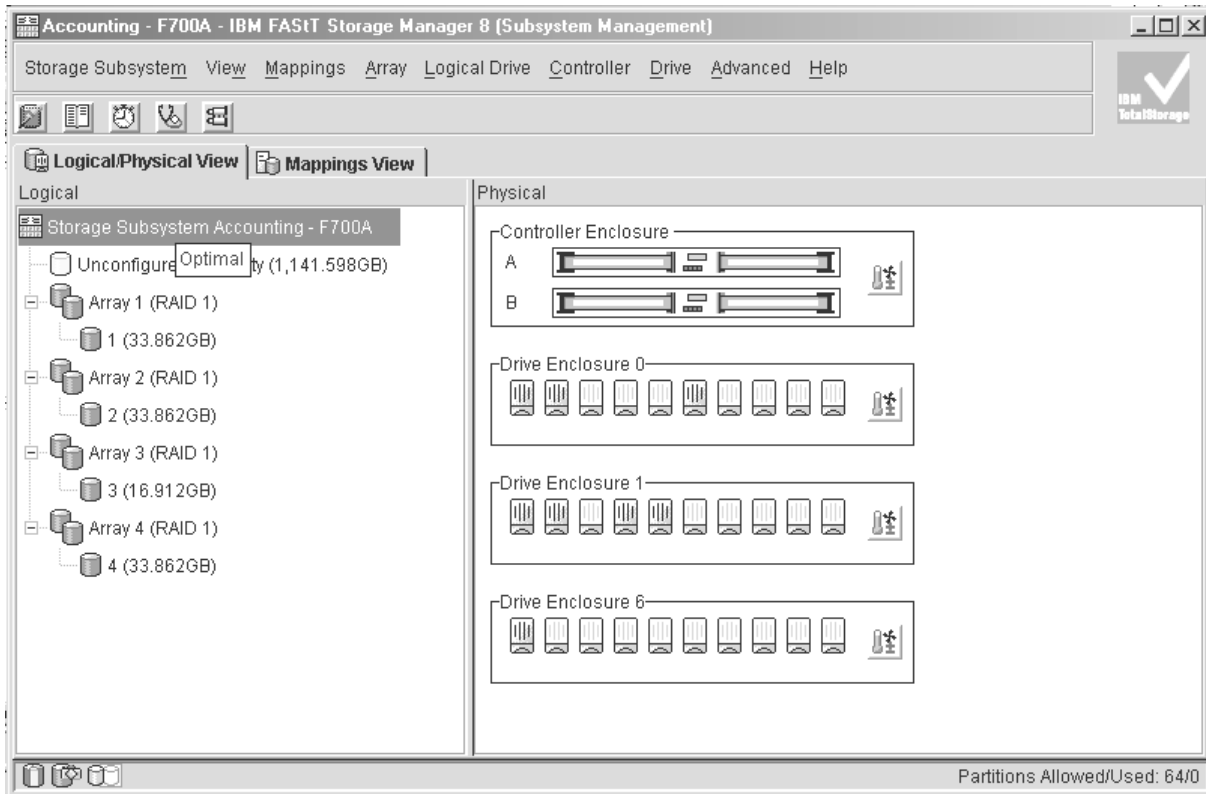
Performing an initial automatic discovery of storage subsystems

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

1. If you are running the SMclient from AIX, first 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.



SJ001148

Figure 12. Enterprise Management window

3. Ensure that all of the attached hosts and storage subsystems are displayed as expected. If they are not, complete the following steps:
 - 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 IBM technical-support representative.
5. Go to “Configuring alert notifications.”

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 FAST Storage Manager Concepts Guide*.

Complete the following steps 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, you can use the storage subsystem profile 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 Storage Partitioning is enabled on your system)

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 19, Chapter 3, “Installing storage management station software on HP-UX systems,” on page 55, or Chapter 4, “Installing storage management station software on Solaris systems,” on page 71 for installation instructions.

Chapter 6. Storage Manager 8.4 with high-availability cluster services

The high-availability clustering services provided by Storage Manager 8.4 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 8.4 with the following cluster services.

Note: AIX clustering is not supported on Storage Manager 8.41.

Operating System	High-Availability Cluster Service
AIX	HACMP ES and ESCRM 4.5, 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 RDAC as your failover driver, and you are using the VERITAS Volume Manager, you must disable Data Multipathing (DMP). Complete one of the following procedures, according to your version of VERITAS Volume Manager.

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`

VERITAS Volume Manager 3.1 or earlier: For VERITAS Volume Manager 3.1 or earlier, you must disable DMP for all devices.

Complete the following steps:

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.

VERITAS Volume Manager 3.2 or later: For VERITAS Volume Manager 3.2 or later, you must disable DMP for FAST storage servers only.

Use option 17 and then option 5 of the submenu of the `vxdiskadm` utility to disable DMP for FAST storage servers. Other devices can use DMP.

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.

Restriction: AIX clustering is not supported on Storage Manager 8.41.

General hardware requirements

For hardware, software, firmware and RDAC requirements, see Appendix A, “AIX system requirements,” on page 119.

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 v8.4**. Accept the license agreement, then click the **v8.4** 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 8.4.

Software requirements

The following HACMP software requirements are in addition to the requirements that are specified in Appendix A, “AIX system requirements,” on page 119:

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

For more information about software requirements, 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 **Important notices and readmes**, then click **HACMP configuration information for FAStT**.

Configuration limitations

The following limitations apply to HACMP configurations:

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

Other HACMP usage notes

The following notations are specific to HACMP environments:

- Concurrent and non-concurrent modes are supported with HACMP 4.4.1 and 4.5 and FASTT running Storage Manager 8.4, including hot standby and mutual takeover.
- HACMP clusters can support from two to 32 servers on each FASTT partition. If you run this kind of environment, be sure to read and understand the AIX device drivers queue depth settings that are described in Chapter 2, “Installing storage management station software on AIX systems,” on page 19.
- You can attach non-clustered AIX hosts to a FASTT that is running Storage Manager 8.4 and is attached to an HACMP 4.5 cluster. However, you must configure the non-clustered AIX hosts on separate host partitions on the FASTT.

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 8.4.

Software requirements

The following combinations of PSSP and GPFS are supported:

Table 17. 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 shows a cluster configuration that contains a single FASTt storage server, with one to four partitions.

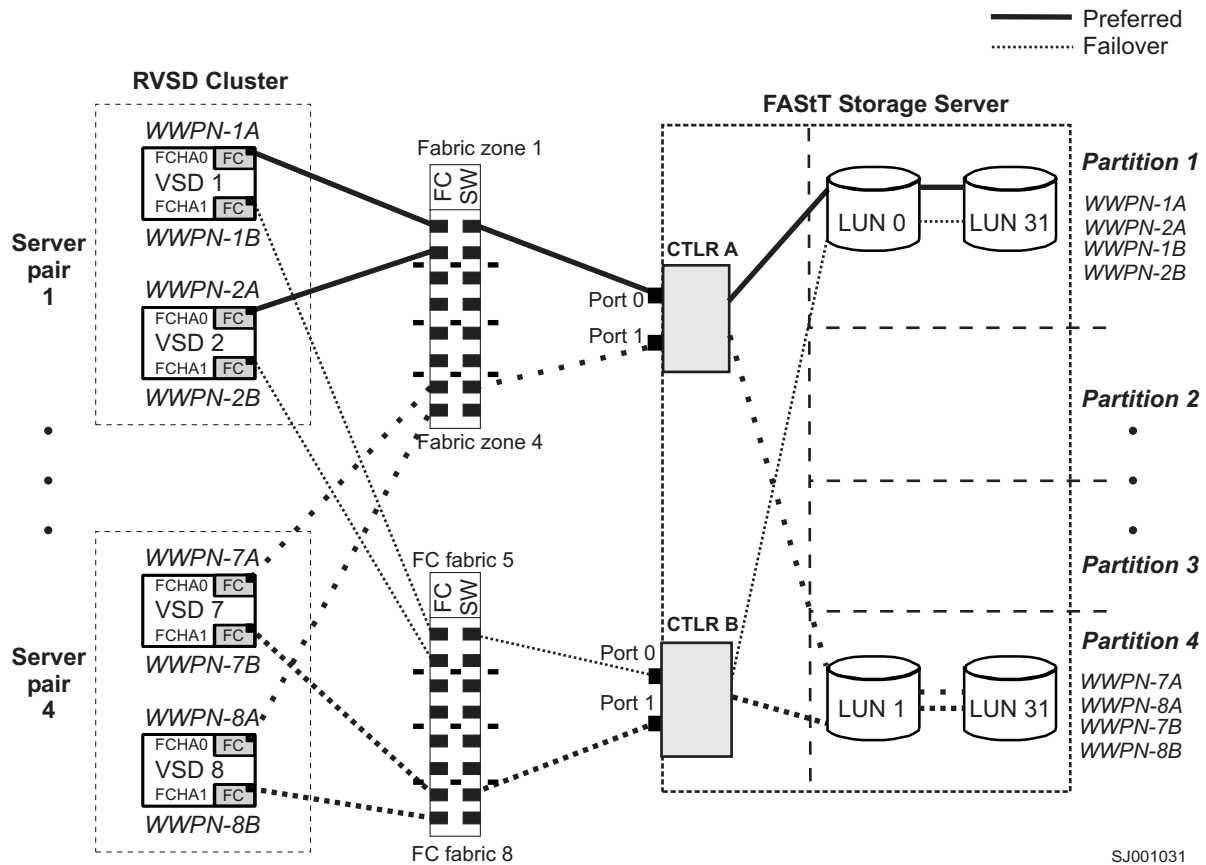


Figure 13. Cluster configuration with single FASTt storage server — one to four partitions

Figure 14 on page 114 shows a cluster configuration that contains three FASTt storage servers, with one partition on each storage server.

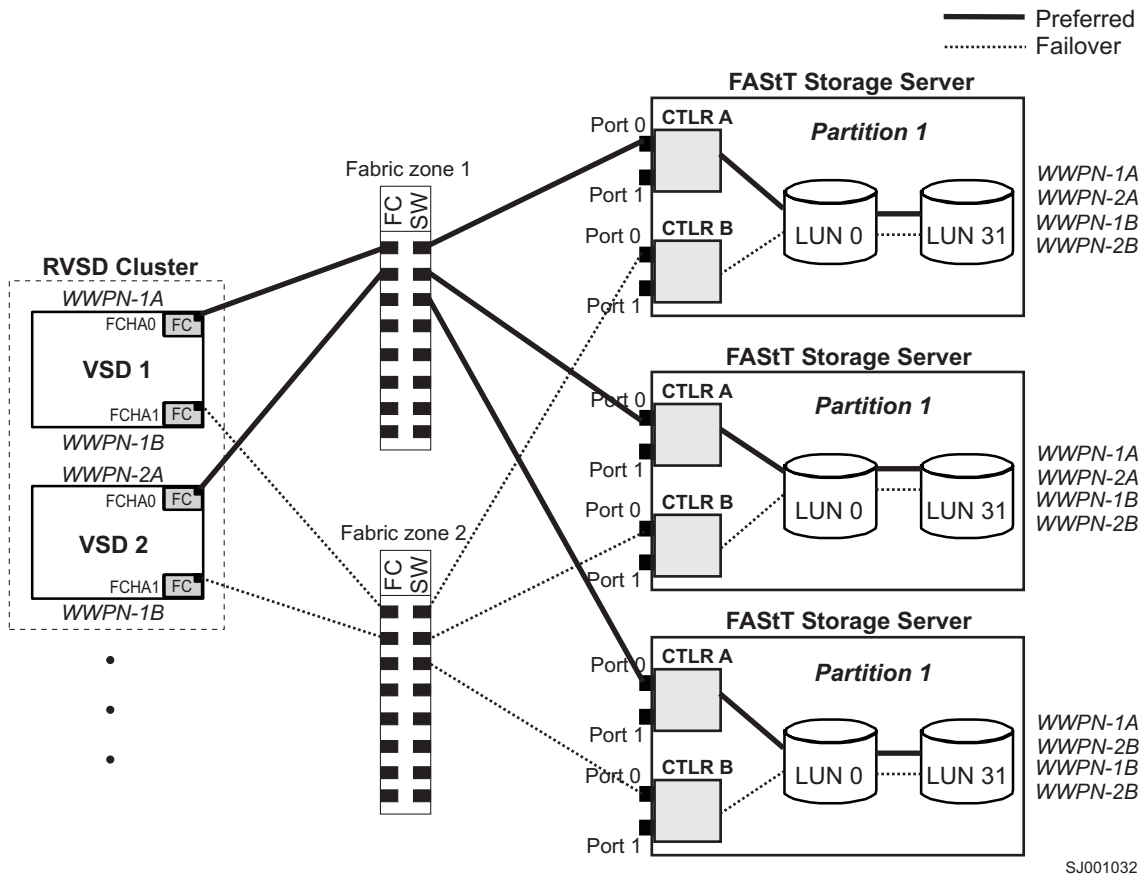


Figure 14. Cluster configuration with three FASTt storage servers — one partition per FASTt

Figure 15 on page 115 shows a cluster configuration that contains four FASTt storage servers, with one partition on each storage server.

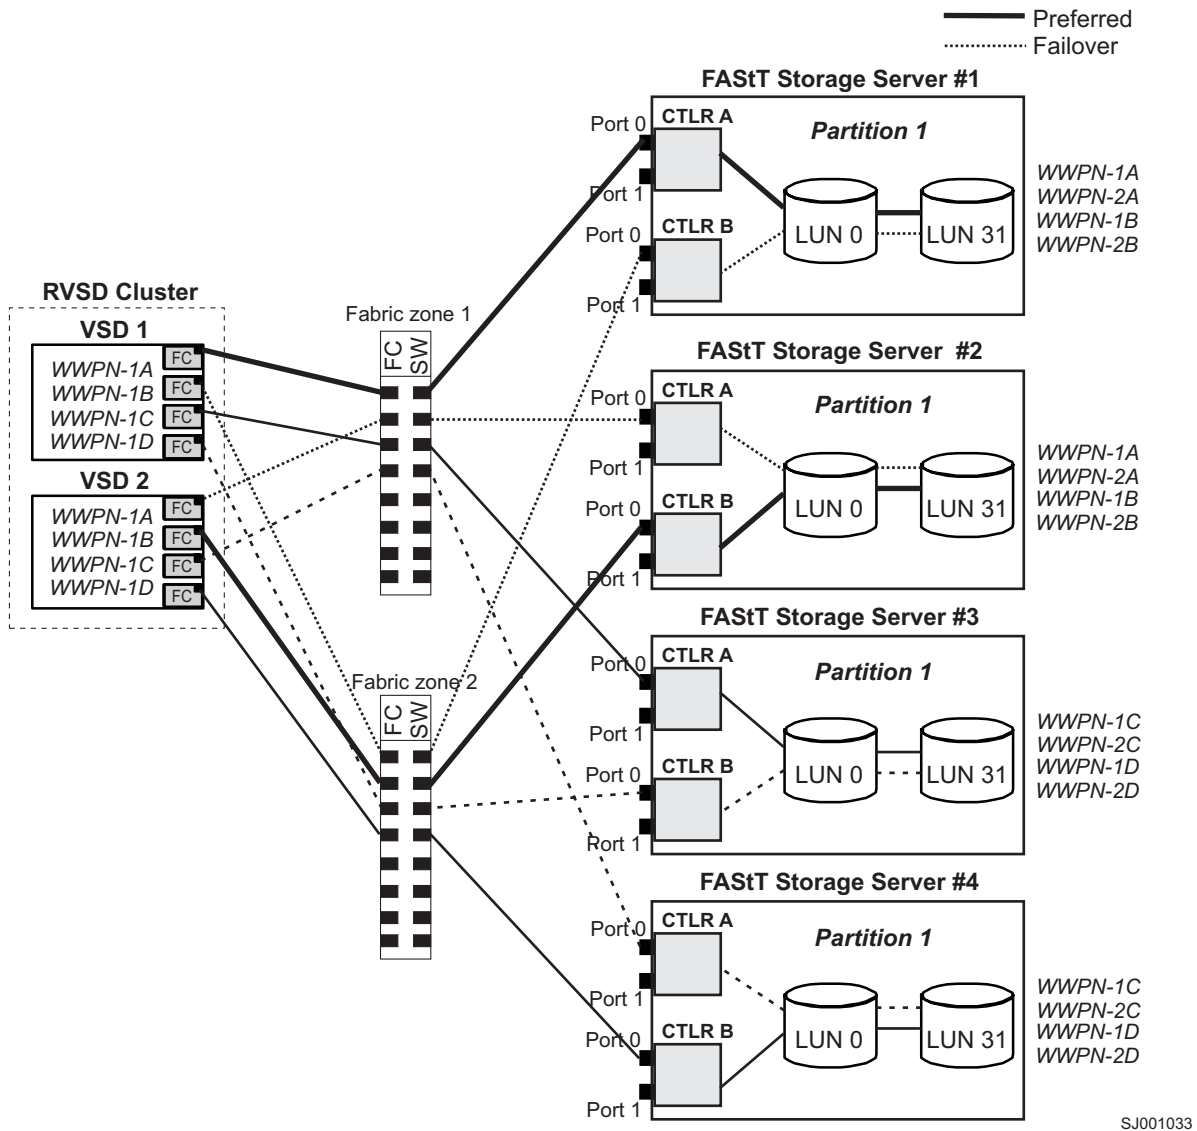


Figure 15. Cluster configuration with four FAST storage servers — one partition per FAST

Figure 16 on page 116 shows a cluster configuration that contains two FAST storage servers, with two partitions on each storage server.

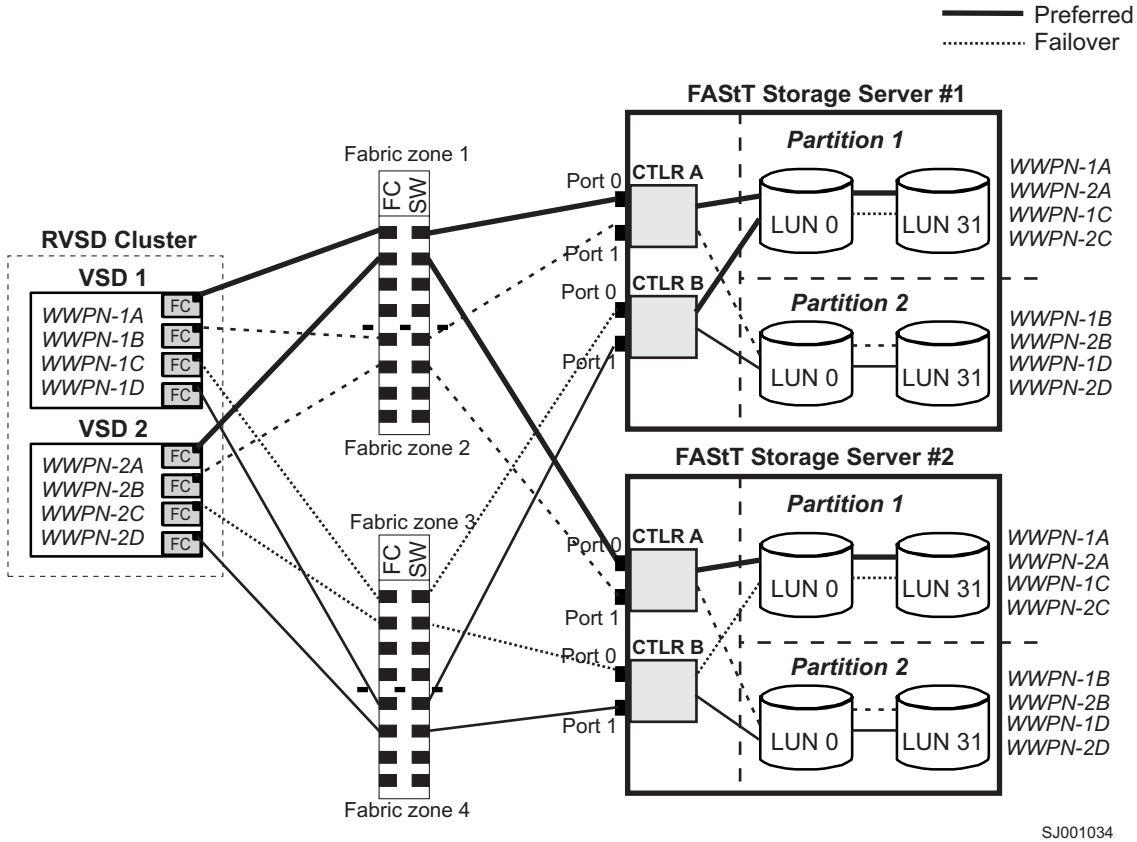


Figure 16. RVSD cluster configuration with two FASTt storage servers — two partitions per FASTt

Figure 17 on page 117 shows an HACMP/GPFS cluster configuration that contains a single FASTt storage server, with one partition.

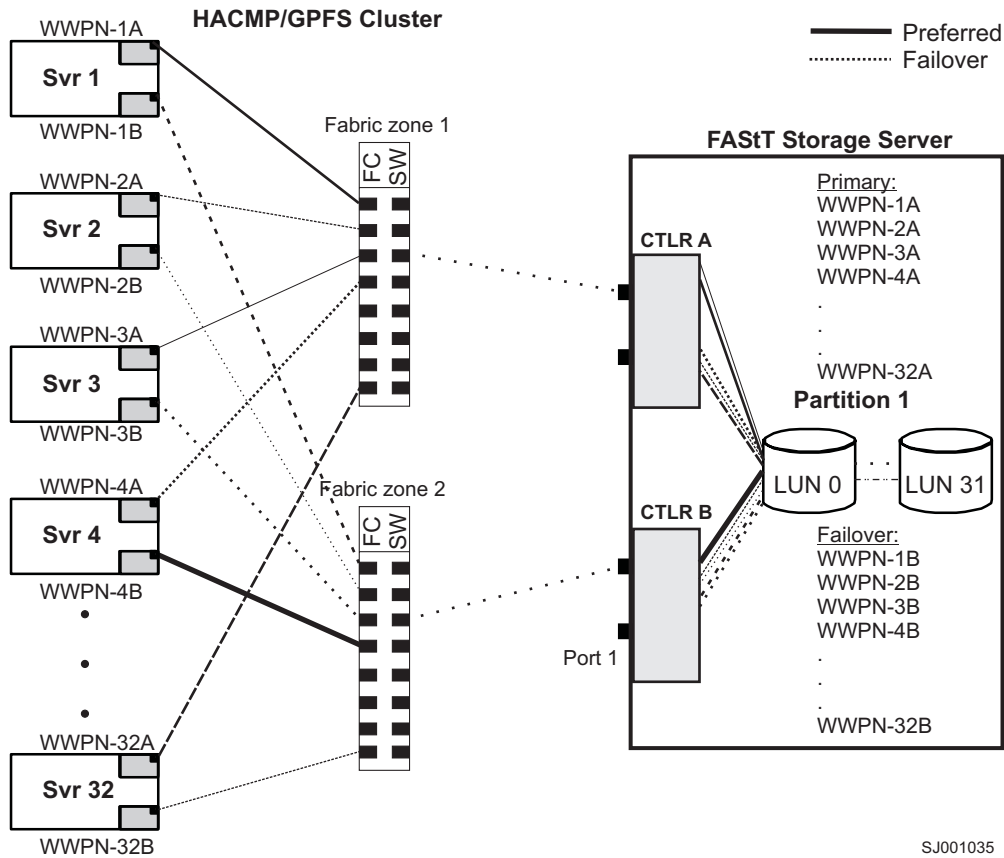


Figure 17. HACMP/GPFS cluster configuration with one FAST storage server — one partition

Figure 18 on page 118 shows an HACMP/GPFS cluster configuration that contains two FAST storage servers, with two partitions on each storage server.

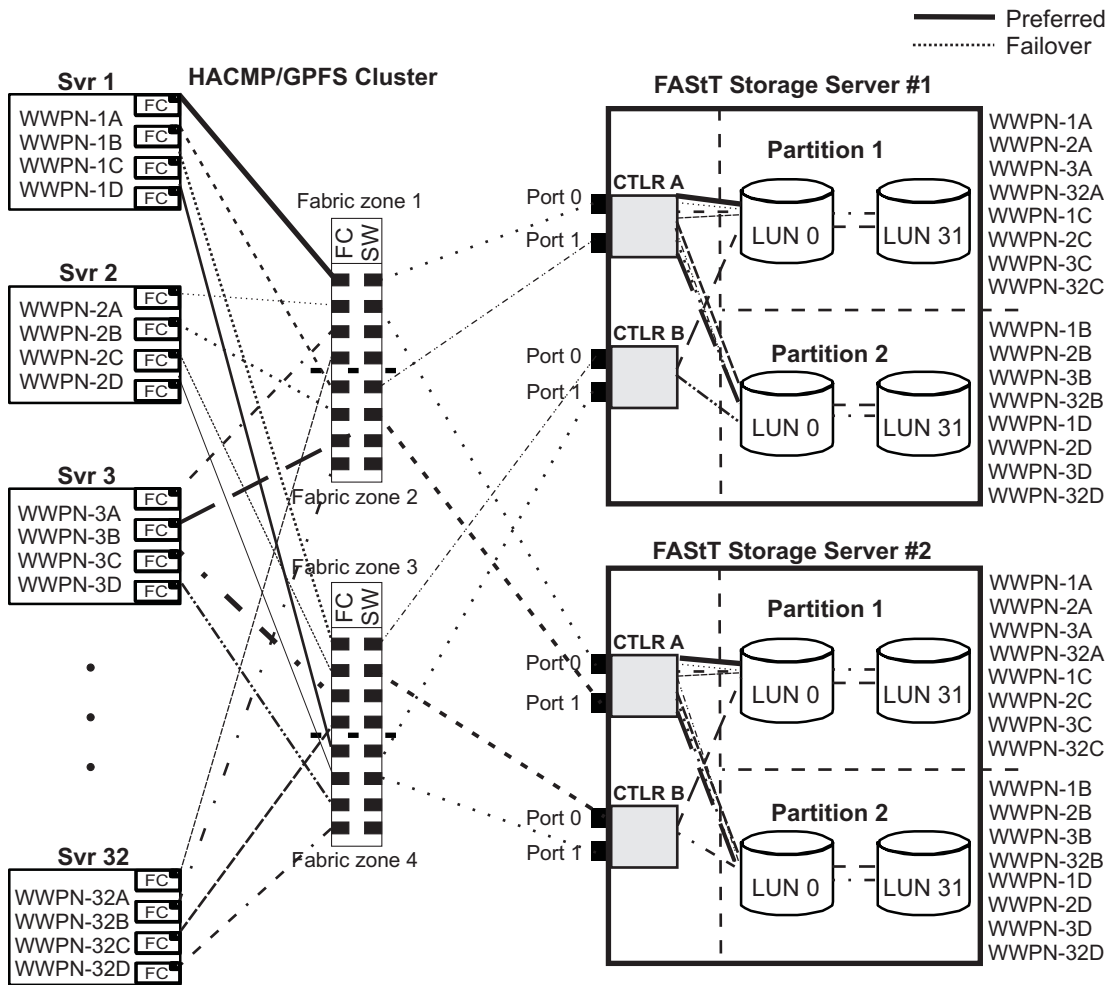


Figure 18. HACMP/GPFS cluster configuration with two FAST storage servers — two partitions per FAST

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 FAST600, FAST700, or FAST900 storage subsystem that is managed by Storage Manager 8.4.

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 FAST 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 FAST storage server. When the page opens, click **Storage Manager v8.4**. Accept the license agreement, then click the **v8.4** link for IBM FAST 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 18, Table 19, or Table 20 on page 120.

RDAC installation requirements

The filesets that are listed in Table 18, Table 19, and Table 20 on page 120 must be the specified versions or later.

Table 18. 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 19. 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 19. 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 20. Filesets required for AIX 5.2 RDAC

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 device driver 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 RDAC on AIX hosts” on page 29.

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 FAST600, FAST700, or FAST900 storage subsystem that is managed by Storage Manager 8.4.

A readme file that contains the latest installation and user information about the storage management software, and hardware components located in the HP-UX FAST 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 FAST storage server. When the page opens, click **Storage Manager v8.4**. Accept the license agreement, then click the **v8.4** link for IBM FAST Storage Manager for HP-UX.

Hardware requirements

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

Software requirements

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 21.

Table 21. 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 FAST600, FAST700, or FAST900 storage subsystem that is managed by Storage Manager 8.4.

A readme file that contains the latest installation and user information about the storage management software, and hardware components located in the Solaris FAST Installation CD. The most recent copy, which 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 FAST storage server. When the page opens, click **Storage Manager v8.4**. Accept the license agreement, then click the **v8.4** link for IBM FAST Storage Manager for Solaris.

Hardware requirements

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

Software requirements

Ensure that the Solaris host is running one of the following operating systems with the appropriate patches. The patches listed in this document can be superseded by more recent versions, so see your Solaris operating system documentation or contact your operating system supplier to ensure that you have the correct patches.

- 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

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
```

where `vg_group_name` is the name of the volume group that is being distributed.

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*.

Complete the following steps 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
```

where `vg_group_name` is the name of the volume group whose paths you want to display.

On node *B*:

```
#vgdisplay -v volume_group_name
```

where `vg_group_name` is the name of the volume group whose paths you want to display.

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 and QLogic host bus adapter settings

The following tables detail settings for the various host bus adapter (HBA) cards for Sun Solaris.

JNI HBA card settings

The JNI cards are not plug-and-play with auto-configuration. Instead, you might need to change the settings or bindings.

Table 22. Configuration settings for FCI-1063

Original value	New value	Comments
fca_nport = 0;	fca_nport = 1;	
ip_disable = 0;	ip_disable = 1;	
failover = 30;	failover = 60;	
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 23. 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 = 60;	Uncomment the line.

Table 23. Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562 (continued)

Original value	New value	Comments
def_hba_binding = JNIC146x	Direct-attached configurations: def_hba_binding = JNIC146x SAN-attached configurations: def_hba_binding = null	Setting depends on whether you have a direct-attached or SAN-attached configuration.
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
```

Table 24. 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:

1. Set `FcPortCfgEnabled = 1;`
2. Restart the host.
3. Set `FcPortCfgEnabled = 0;`
4. 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 25. Configuration settings for FC64-1063

Original value	New value	Comments
<code>fca_nport = 0;</code>	<code>fca_nport = 1;</code>	
<code>ip_disable = 0;</code>	<code>ip_disable=1;</code>	
<code>failover = 0;</code>	<code>failover =30;</code>	
<code>busy_retry_delay = 5000;</code>	<code>busy_retry_delay = 5000;</code>	
<code>link_recovery_delay = 1000;</code>	<code>link_recovery_delay = 1000;</code>	
<code>scsi_probe_delay = 5000;</code>	<code>scsi_probe_delay = 5000;</code>	
<code>def_hba_binding = "fcaw*";</code>	Direct attached configurations: <code>def_hba_binding = "fcaw*";</code> SAN-attached configurations: <code>def_hba_binding = "nonJNI";</code>	Setting depends on whether you have a direct-attached or SAN-attached configuration.
<code>def_wwnn_binding = "\$xxxxxx"</code>	<code>def_wwnn_bindindef_hba_binding = "nonjni"; g = "xxxxxx"</code>	Uncomment the line.
<code>def_wwnn_binding = "\$xxxxxx"</code>	Same as the original entry.	Uncomment the line.
Will be added by reconfigure script	<code>name="fcaw" parent="<physical path>"unit-address="<#>"</code>	
Will be added by reconfigure script	<code>target0_hba="fcaw0"</code> <code>target0_wwpn="<controller wwpn>;"</code>	
Will be added by reconfigure script	<code>name="fcaw" parent="<physical path>"unit-address="<#>"</code>	
Will be added by reconfigure script	<code>target0_hba="fcaw0"</code> <code>target0_wwpn="<controller wwpn>;"</code>	

Note: You might need to run the `/etc/raid/bin/genscsiconf` reconfigure script from the shell prompt.

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

QLogic HBA card settings

The QLogic cards are not plug-and-play with auto-configuration. Instead, you might need to change the settings or bindings.

Table 26. Configuration settings for QL2342

Original value	New value	Comments
<code>hba0-max-frame-length=2048;</code>	<code>hba0-max-frame-length=2048</code>	Use the default.
<code>hba0-execution-throttle=16;</code>	<code>hba0-execution-throttle=255;</code>	Change.
<code>hba0-login-retry-count=8;</code>	<code>hba0-login-retry-count=30;</code>	Change.

Table 26. Configuration settings for QL2342 (continued)

Original value	New value	Comments
hba0-enable-adapter-hard-loop-ID=0;	hba0-enable-adapter-hard-loop-ID=1;	Change.
hba0-adapter-hard-loop-ID=0;	hba0-adapter-hard-loop-ID=0;	Needs to be a unique number.
hba0-enable-LIP-reset=0;	hba0-enable-LIP-reset=0;	Use the default.
hba0-enable-LIP-full-login=1;	hba0-enable-LIP-full-login=1;	Use the default.
hba0-enable-target-reset=0;	hba0-enable-target-reset=0;	Use the default.
hba0-reset-delay=5	hba0-reset-delay=8	Change.
hba0-port-down-retry-count=30;	hba0-port-down-retry-count=70;	Change.
hba0-maximum-luns-per-target=8;	hba0-maximum-luns-per-target=0;	Change.
hba0-connection-options=2;	hba0-connection-options=2;	Use the default.
hba0-fc-tape=1;	hba0-fc-tape=0;	Change.

Appendix F. Additional FAStT documentation

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

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

FAStT900 Fibre Channel Storage Server library

Table 27 associates each document in the FAStT900 Fibre Channel Storage Server library with its related common user tasks.

Table 27. TotalStorage FAStT900 Fibre Channel Storage Server document titles by user tasks

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

Table 27. TotalStorage FAStT900 Fibre Channel Storage Server document titles by user tasks (continued)

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

FAStT700 Fibre Channel Storage Server library

Table 28 associates each document in the FAStT700 Fibre Channel Storage Server library with its related common user tasks.

Table 28. TotalStorage FAStT700 Fibre Channel Storage Server document titles by user tasks

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

FAStT600 Fibre Channel Storage Server library

Table 29 associates each document in the FAStT600 Fibre Channel Storage Server library with its related common user tasks.

Table 29. TotalStorage FAStT600 Fibre Channel Storage Server document titles by user tasks

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

FASTt-related documents

Table 30 associates each of the following documents related to FASTt operations with its related common user tasks.

Table 30. 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 TotalStorage FASTt EXP100 Storage Expansion Unit Release Notes, GC26-7619	✓	✓				
IBM TotalStorage FASTt EXP100 Storage Expansion Unit Installation and Users Guide, GC26-7601	✓	✓		✓	✓	✓
Fibre Channel Solutions - IBM FASTt EXP500 Installation and User's Guide, 59p5637	✓	✓		✓	✓	✓
IBM EXP700 Storage Expansion Unit Installation and User's Guide, 32P0178	✓	✓		✓	✓	✓
IBM Netfinity® Fibre Channel Cabling Instructions, 19K0906		✓				
IBM Fibre Channel SAN Configuration Setup Guide, 25P2509	✓		✓	✓	✓	

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

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.

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

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

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 LUN 60, 77
- access volume 2
- access volumes 12
- address
 - for e-mail comments xvi
 - for reader comment form xvii
 - of IBM xvii
 - of the IBM director of licensing 137
- agent software
 - AIX 23
 - HP-UX 58
 - package 8
 - Solaris 76
- AIX
 - client software 22
 - configuration of storage subsystems 24, 30
 - direct-attached configuration 28
 - fget_config command 42
 - firmware
 - requirements 19
 - upgrading 25
 - hardware requirements 119
 - identifying
 - controller ID numbers 39
 - device names and bus numbers 40
 - FlashCopy volume types 48
 - logical drives by operating system device names 46
 - lsattr command 45
 - lsdev command 41
 - NVSRAM
 - requirements 19
 - upgrading 25
 - operating system requirements 16, 17, 19
 - RDAC driver
 - configuring 30
 - installing 29, 30
 - requirements 29, 119
 - verifying 29
 - requirements
 - client software 22
 - firmware 19
 - hardware 19, 119
 - operating system 16, 17
 - RDAC driver 29, 119
 - software 119
 - system 119
 - restrictions 21
 - SAN-attached configuration 28
 - SMagent 23
 - SMclient 23
 - SMruntime 23
 - software requirements 119
 - system requirements 119
 - volumes, redistributing in case of failure 54

- alert notifications, configuring 107
- attributes
 - dac 32, 36
 - dar 31, 36
 - definitions 31
 - hdisk 33, 36, 37
 - LUN 33, 36, 37
 - RDAC driver 30
- audience of this document xi
- Auto Volume Transfer (AVT) 69, 103

B

- bus numbers and device names, identifying 40

C

- cabling instructions xi
- cache mirroring 37, 54, 112
- capture state information feature 107
- catastrophic failure 107
- client software
 - AIX 22, 23
 - HP-UX 56, 57
 - installing 56
 - package 8, 9
 - Solaris 73, 74
- cluster services
 - AIX requirements 111, 112
 - configurations 105
 - HACMP ES and ESCRM 109
 - hardware requirements 109
 - HP-UX requirements 109
 - MC/Service Guard 109
 - PSSP with GPFS 109
 - Solaris requirements 109, 111
 - system dependencies 109
 - VERITAS Cluster Server 109, 110
 - VERITAS Volume Manager 109, 110
- commands
 - fget_config (AIX) 42
 - lsattr (AIX) 45
 - lsdev (AIX) 41
- comments about this document, how to send xvi
- complete SM SW installation xi
- Concepts Guide xi, xiii
- configuration details, MC Service Guard 125
- configuration recovery 107
- configure storage hardware xi
- configure storage subsystems on host xi
- configuring
 - alert notifications 107
 - direct-attached 28, 63, 79
 - RDAC driver 30, 110
 - SAN-attached 28, 63, 79
 - Solaris with JNI host bus adapter cards 80, 81
 - Solaris with QLogic host bus adapter cards 80

- configuring (*continued*)
 - storage subsystems
 - AIX 24, 28
 - HP-UX 60, 63
 - initial automatic discovery 105
 - Solaris 77, 79
- connect power xi
- controller state capture 107
- controllers
 - firmware 1, 29, 60, 78
 - ID numbers, identifying 39
 - IP addresses 17
- Copy Services Guide xiii

D

- dac
 - and fget_config command 42
 - and lsattr command 45
 - and RDAC 30
 - attributes 32, 36
 - identifying controller ID numbers 39
 - showing status and location 41
- dar
 - and fget_config command 42
 - and lsattr command 45
 - and RDAC 30
 - attributes 31, 36
 - showing status 41
 - two dars on one subsystem, causes 30
- DCE (dynamic capacity expansion) 51
- delayed_fail 49
- determine management method xi
- devices
 - identification 38
 - names and bus numbers, identifying 40
- direct (out-of-band) management method 12
- direct-attached configuration 28, 63, 79
- disk configuration information 107
- documentation
 - documents 1
 - FASTT xii, 131
 - FASTT Storage Manager Version 8.4 xiii
 - files 1
 - online help xv, 2, 106, 108
 - related 135
 - Web sites xvi
- DVE (dynamic volume expansion) 51
- dynamic capacity expansion (DCE) 51
- Dynamic Multipathing (DMP)
 - configuring VERITAS Volume Manager for 98
 - description 9
 - disabling 110
 - installing 92
 - installing the SMibmasl package for 97
 - installing VERITAS Volume Manager packages for 94
 - preparing for installation 92
 - system requirements 92
- dynamic tracking 50
- dynamic volume expansion (DVE) 51

E

- e-mail address xvi, 107
- edition notice ii
- enabling multipath I/O with PV-links 64
- Enterprise Management window
 - description 8
 - online help 2
- EXP100
 - hardware and firmware requirements 19
 - release notes 2
 - restrictions 2, 19, 21
- EXP500
 - hardware and firmware requirements 19, 55, 71
- EXP700 7
 - hardware and firmware requirements 19, 55, 71

F

- failure support
 - cluster services 109
 - DMP driver 9
 - RDAC driver 9
 - redistributing volumes 54, 69, 103
- fast I/O failure 49
- fast_fail 49
- FASTT documentation xii, 131
- FASTT Storage Manager
 - related documents 135
 - Version 8.4 Library xiii
- FASTT Storage Manager 8.4
 - user tasks xiii
- FASTT, introduction to 17
- FASTT600
 - restrictions 5, 7
- FASTT600 Fibre Channel Storage Server library 134
- FASTT600 Turbo 5
- FASTT700 Storage Server library 133
- FASTT900 Fibre Channel Storage Server library 131
- FC cabling instructions xi
- FC HW maintenance and problem determination guide xi
- FC planning and integration user's guide and service information xi
- FC storage server installation guide xi
- fcsci attributes 49, 50
- fget_config command 42
- filesets 119
- firmware
 - AIX
 - requirements 19
 - upgrading 25
 - verifying 29
 - description 1
 - downloading latest versions 25, 61, 78
 - HP-UX
 - requirements 55
 - upgrading 61
 - Solaris
 - requirements 71
 - upgrading 78

firmware (*continued*)
version 5.41.xx.xx, warning 19

G

GPFS with PSSP 109

H

HACMP ES and ESCRM 109
hardware installation process xi
hardware maintenance and problem determination xi
hardware requirements
AIX 19, 119
HP-UX 55, 121
Solaris 71, 123
HBA
binding to storage controllers 82, 84, 85
modifying settings 81, 83
selecting host port
AIX 26
setting host port
HP-UX 62
Solaris 88
HBA installation and user's guide xi
hdisk
and fget_config command 42
and lsattr command 45
and RDAC 30
attributes 33, 36, 37
determining preferred path 42
setting queue depth 37
showing status and location 41
help xv, 2, 106, 108
heterogeneous environment 105
high-availability cluster services 109
host bus adapter installation and user's guide xi
host group, setting up
AIX 26
HP-UX 62
Solaris 88
host software
AIX 23
HP-UX 58, 59
packages 9
Solaris 75, 76
host-agent (in-band) management method 10
how to send your comments xvi
HP-UX
client software 56
cluster services, requirements 109
configuration of storage subsystems 60
direct-attached configuration 63
firmware
requirements 55
upgrading 61
hardware requirements 55, 121
host software
installing 58
requirements 58

HP-UX (*continued*)

NVSRAM
requirements 55
upgrading 61
operating system requirements 16, 17
requirements
client software 56
cluster services 109
firmware 55
hardware 55, 121
host software 58
NVSRAM 55
operating system 16, 17
software 121
system 121
SAN-attached configuration 63
SMagent 58
SMclient 57
SMruntime 57
SMutil 59
software requirements 121
system requirements 121
volumes, redistributing in case of failure 69
HyperTerminal version 17

I

IBM
address xvii
director of licensing address 137
e-mail address xvi, 107
identifying
AIX
controller ID numbers 39
device names and bus numbers 40
logical drives by operating system device names 46
in-band (host-agent) management method 10
install and verify SM SW on host and workstation xi
install network hardware xi
install storage expansion unit xi
install storage server/RAID enclosures in rack xi
installation and support OS guides xi
installation guide, storage server xi
installation planning xi
installation process xi
installing
AIX
client software 22, 23
RDAC driver 29
client software
AIX 22
HP-UX 56
SMclient 23, 57, 74
SMruntime 23, 57, 73
Solaris 73
host software
HP-UX 58
SMagent 23, 58, 76
SMutil 59, 76
Solaris 75

installing (*continued*)
 HP-UX
 client software 56, 57
 host software 58
 SMutil 59
 process diagram for 14
 sequence of 9
 Solaris
 client software 73, 74
 host software 75, 76
 RDAC driver 90
IP addresses for FASiT controllers 17

J

JNI
 cards 80, 81
 installing driver 80
 modifying HBA settings 81, 83
 settings 127

L

loop settings 81, 83
lsattr command 45
lsdev command 41
LUNs
 and access volumes 12
 and fget_config command 42
 attributes 33, 36, 37
 mapping
 AIX 27
 HP-UX 63
 Solaris 89
 maximum number per partition 3

M

make FC connections xi
management method
 direct (out-of-band) 12
 host-agent (in-band) 10
management station 1
 See storage management station
MC Service Guard 125
MC/Service Guard 109
multipath I/O with PV-links 64

N

notice, edition ii
notices
 general 137
NVSRAM
 AIX
 requirements 19, 21
 downloading latest versions 25, 61, 78
 HP-UX
 requirements 55
 upgrading 61

NVSRAM (*continued*)
 Solaris
 requirements 71
 upgrading 78

O

online help xi, xv, 2, 106, 108
operating system
 AIX 16, 17, 19, 119
 HP-UX 16, 17, 121
 requirements 55
 requirements
 AIX 19
 HP-UX 55
 Solaris 71
 Solaris 16, 17, 123
 requirements 71
out-of-band (direct) management method 12
overview of heterogeneous hosts 105

P

premium features
 copy services 5
 enabling 6
 storage partitioning 6
 host group 26, 62, 88
prepare for SM SW installation xi
prepare network connection xi
prerequisites
 AIX
 client software 22
 firmware 19
 hardware 19, 119
 operating system 16, 17
 RDAC driver 29, 119
 software 119
 system 119
 cluster services 109
 HP-UX
 client software 56
 cluster services 109
 firmware 55
 hardware 55, 121
 host software 58
 NVSRAM 55
 operating system 16, 17
 software 121
 system 121
 Solaris
 client software 73
 cluster services 109, 111
 firmware 71
 hardware 71, 123
 host software 75
 NVSRAM 71
 operating system 16, 17
 RDAC driver 75
 software 123
 system 123

products, developed 137
PSSP with GPFS 109
PV-links 64

Q

QLogic
 cards 80
 settings 129
queue depth for hdisk devices 37

R

RAID controller enclosure unit installation guide xi
RDAC driver
 AIX
 configuring 30
 installing 29, 30
 requirements 29, 119
 attributes 30
 description 9
 IDs 110
 installing 89
 Solaris
 installing 89, 90
 requirements 75
reader comment form xvii
readme file 1
recover configuration 107
redistributing volumes in case of failure 54, 69, 103
removing previous versions of Storage Manager
 AIX 22
renaming storage subsystem 107
requirements
 AIX
 client software 22
 firmware 19
 hardware 19, 119
 operating system 16, 17
 RDAC driver 29, 119
 software 119
 system 119
 cluster services 109
 HP-UX
 client software 56
 cluster services 109
 firmware 55
 hardware 55, 121
 host software 58
 NVS RAM 55
 operating system 16, 17
 software 121
 system 121
 Solaris
 client software 73
 cluster services 109, 111
 firmware 71
 hardware 71, 123
 host software 75
 NVS RAM 71
 operating system 16, 17

requirements (*continued*)
 Solaris (*continued*)
 RDAC driver 75
 software 123
 system 123
resources
 documents 1
 e-mail address xvi, 107
 files 1
 online help xv, 2, 106, 108
 Web sites xvi, 6
restrictions
 AIX 21
 FAST600 5, 7
 in-band management 10
runtime software
 package 7

S

SAN-attached configuration 28, 63, 79
SATA
 hardware and firmware requirements 19
 release notes 2
 restrictions 2, 19, 21
Script Commands xiii
sending your comments to IBM xvi
services offered in the U.S.A. 137
set link speed xi
setting
 IP addresses 17
 NMS for SNMP notification 107
SM concepts guide xi
SM installation and support guide for Linux xi
SM installation and support guide for Novell
 NetWare xi
SM installation and support guide for Windows 2000 xi
SM installation and support guide for Windows NT xi
SMagent
 AIX installing 23
 description 8
 HP-UX installing 58
 installing 76
 Solaris installing 76
SMclient
 AIX installing 23
 description 8
 HP-UX installing 57
 Solaris installing 74
SMruntime
 AIX installing 23
 description 7
 HP-UX installing 57
 Solaris installing 73
SMutil
 description 9
 HP-UX installing 59
 Solaris installing 76
SNMP notification 107
software installation process xi

- software package
 - RDAC 9
 - SMagent 8
 - SMclient 8
 - SMruntime 7
 - SMutil 9
 - software requirements
 - AIX 119
 - HP-UX 121
 - Solaris 123
 - Solaris
 - client software 73
 - cluster services requirements 109, 111
 - configuration of storage subsystem 77
 - direct-attached configuration 79
 - DMP 92, 94, 97, 98, 101
 - firmware
 - requirements 71
 - upgrading 78
 - hardware requirements 71, 123
 - host software 75
 - JNI host bus adapter cards 80, 81
 - NVSRAM
 - requirements 71
 - upgrading 78
 - operating system requirements 16, 17
 - QLogic host bus adapter cards 80
 - RDAC driver
 - installing 90
 - requirements 75
 - requirements
 - client software 73
 - cluster services 109
 - firmware 71
 - hardware 71, 123
 - host software 75
 - NVSRAM 71
 - operating system 16, 17
 - RDAC driver 75
 - software 123
 - system 123
 - SAN-attached configuration 79
 - SMagent
 - installing 76
 - SMclient
 - installing 74
 - SMruntime
 - installing 73
 - SMutil
 - installing 76
 - software requirements 123
 - system requirements 123
 - VERITAS 92, 94, 98, 101
 - volumes, redistributing in case of failure 103
 - starpubs xvi
 - start server xi
 - storage expansion unit installation guide xi
 - storage management station 1
 - Storage Manager software
 - installation sequence 9
 - introduction 1
 - Storage Manager software (*continued*)
 - new features 2
 - software packages 7
 - version 8.41 xi, 2, 5, 19, 23, 109, 111
 - where to obtain 1
 - storage server installation guide xi
 - storage subsystem
 - adding to SMclient
 - AIX 24
 - HP-UX 60
 - Solaris 78
 - cluster services 109
 - configuring
 - AIX 24
 - HP-UX 60
 - procedures 105
 - Solaris 77
 - initial automatic discovery 105
 - introduction 1
 - logical drives
 - identifying by operating system device names 46
 - management methods
 - direct (out-of-band) 12
 - host-agent (in-band) 10
 - management tasks 108
 - renaming 107
 - storage subsystem profile, saving 107
 - Subsystem Management window
 - description 8
 - online help 2
 - switch
 - in a SAN-attached configuration 28, 63, 79
 - modifying JNI HBA settings 81, 83
 - zoning 28, 63, 79
 - system requirements
 - AIX 119
 - HP-UX 121
 - Solaris 123
- ## T
- tasks by documentation title xii, 131
 - trademarks 137
- ## U
- universal transport mechanism 30
 - upgrading
 - AIX
 - NVSRAM 25
 - HP-UX
 - firmware 61
 - NVSRAM 61
 - Solaris
 - firmware 78
 - NVSRAM 78
 - user tasks
 - FASTT Storage Manager xiii
 - FASTT600 Fibre Channel Storage Server 134
 - FASTT700 Fibre Channel Storage Server 133
 - FASTT900 Fibre Channel Storage Server 131

user tasks *(continued)*
 related documents 135
utility software
 package 9
utm device 30

V

variable settings for JNI adapter cards 81
verify server operation w/ LEDs xi
VERITAS
 Cluster Server 109, 110
 Dynamic Multipathing (DMP) 9, 92, 94, 98, 109,
 110
 Enterprise Administrator 101
 File System 92, 98
 Volume Manager 9, 92, 94, 98, 109, 110
volumes
 redistributing in case of failure 54, 69, 103

W

Web sites xvi, 6, 21, 29, 111

Z

zoning switches 28, 63, 79

Readers' Comments — We'd Like to Hear from You

IBM TotalStorage FAStT Storage Manager Version 8.4x
Installation and Support Guide for AIX, HP-UX, and Solaris

Publication No. GC26-7622-01

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

GC26-7622-01



Spine information:



IBM TotalStorage FAS*t*T Storage
Manager Version 8.4x

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