

IBM TotalStorage FASSt Storage Manager Version
8.21 for
UNIX and AIX Environments

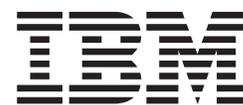


Installation and Support Guide

Read Before Using

The IBM Agreement for Licensed Internal Code is included in this book. Carefully read the agreement. By using this Machine or feature/MES you agree to abide by the terms of this agreement and applicable copyright laws.

IBM TotalStorage FAStT Storage Manager Version
8.21 for
UNIX and AIX Environments



Installation and Support Guide

Note:

Before using this information and the product it supports, read the information in "Notices" on page 135.

Second Edition (November 2002)

This edition replaces GC26-7490-00

© Copyright International Business Machines Corporation 2002. 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
How this document is organized	xi
Related publications	xii
Documents	xii
Additional reference	xiii
Online help	xiii
Web sites	xiii
How to send your comments.	xiv
Chapter 1. Introduction	1
Storage Manager 8.21 features	2
Storage Manager 8 software packages	3
Storage Manager 8 client software package.	4
Storage Manager 8 agent software package	5
Storage Manager 8 utility software package.	5
RDAC	5
Storage subsystem management methods	5
Host-agent (in-band) management method	6
Direct (out-of-band) management method	7
Operating system requirements	9
Setting up IP addresses for FASiT storage controllers	9
Chapter 2. Installing storage management station software on AIX systems 11	
Hardware and firmware requirements.	11
Creating a direct-attached configuration	12
Creating a SAN-attached configuration	12
AIX restrictions	13
Installing the host software on AIX hosts	13
Prerequisites.	13
Installing SMruntime on AIX hosts	13
Installing SMclient on AIX hosts.	14
Installing host RDAC software on AIX hosts	14
Prerequisites.	14
Installing RDAC on AIX hosts	15
Performing the initial configuration of storage subsystems on AIX hosts	15
Updating FASiT firmware and NVSRAM	16
Setting up an AIX host group.	17
AIX configuration information.	17
Viewing and setting attributes of the RDAC driver for AIX	17
Viewing ODM attributes in AIX	22
Changing ODM attribute settings in AIX	23
Verifying the installation and configuration of AIX hosts	24
Identifying the controller ID numbers	25
Identifying device names and bus numbers	26
Identifying logical drives by operating system device names	29
Disk array errors	31
Redistributing volumes in case of failure	32

Chapter 3. Installing storage management station software on HP-UX systems	35
Hardware and firmware requirements	35
Installing the host software on HP-UX hosts	36
Prerequisites.	36
Installing SMruntime on HP-UX hosts	36
Installing the SMclient on HP-UX hosts	37
Installing the client software on HP-UX hosts	38
Installing the SMagent on HP-UX hosts	38
Installing SMutil on HP-UX hosts	38
Configuring storage and partitioning for heterogeneous environments	39
Enabling multipath I/O with PV-links	41
Identifying the controller ID numbers	44
Identifying device names and bus numbers	44
Chapter 4. Installing storage management station software on Solaris systems	47
Hardware and firmware requirements	47
Installing the runtime software on Solaris hosts	48
Prerequisites.	48
Procedure.	48
Installing the client software on Solaris hosts	49
Prerequisites.	49
Installing SMclient on Solaris hosts	49
Installing host software on Solaris hosts.	50
Prerequisites.	51
Installing RDAC on Solaris hosts	51
Installing SMagent on Solaris hosts	52
Installing SMutil on Solaris Host	53
Performing the initial configuration of storage subsystems on Solaris hosts	54
Uninstalling Storage Manager 8.21 software	55
Default partitioning for Solaris devices	55
Configuring JNI host bus adapter cards	56
Variable settings for JNI adapter cards	56
Setting the Solaris system host type	57
Chapter 5. Completing the software installation.	59
Failover protection	59
Changing the host type	59
Heterogeneous hosts overview	60
Configuring storage subsystems	60
Starting the Enterprise Management window	60
Configuring alert notifications.	62
Starting the Subsystem Management window	63
Renaming storage subsystems	64
Performing optional storage subsystem management tasks	64
Chapter 6. Storage Manager 8.21 with high-availability cluster services	67
General information	67
Prerequisites for HP-UX	67
Prerequisites for Solaris and Veritas Cluster Server	67
General hardware requirements.	67
System dependencies	67
Storage Manager 8.21 with General Parallel File System (GPFS)	68
Minimum AIX, PSSP, HACMP, and GPFS software requirements	68
FAST software and firmware requirements	69

Configuration limitations	69
Other usage notes specific to PSSP/GPFS environments	69
GPFS/PSSP logical configuration diagrams	69
Chapter 7. FlashCopy and the FASiT Storage Server.	77
Overview of FlashCopy	77
Enabling FlashCopy	78
Obtaining the feature enable identifier	78
Generating the feature key file	78
Using the feature key file to enable FlashCopy	79
FlashCopy parameters	79
FlashCopy logical drives	80
Viewing FlashCopy logical drive failure settings	80
Estimating FlashCopy repository logical drive capacity	83
Estimating FlashCopy repository life	84
Using FlashCopy	85
Checking the status of the FlashCopy premium feature	86
Creating a FlashCopy logical drive.	87
Mapping a FlashCopy drive to a host.	98
Viewing the FlashCopy logical drive status	101
Disabling a FlashCopy logical drive	106
Recreating a FlashCopy logical drive	108
Resizing a FlashCopy repository drive	109
Deleting a FlashCopy drive	111
Using FlashCopy with UNIX regular disks.	113
Instructions for Solaris.	113
Instructions for HP-UX: Logical Drive Manager	116
Script editor and command-line interface	120
Missing logical drives	120
Appendix A. AIX system requirements	123
Hardware requirements	123
Software requirements	123
RDAC installation requirements	123
Appendix B. HP-UX system requirements	125
Appendix C. Solaris system requirements	127
Appendix D. MC Service Guard configuration details	129
Appendix E. JNI Host Bus Adapter settings	131
Notices	135
Trademarks.	135
IBM agreement for licensed internal code.	136
Actions you must not take	137
Glossary	139
Index	141

Figures

1. Host-agent (in-band) managed storage subsystems	6
2. Direct (out-of-band) managed storage subsystems	8
3. Controller Properties window	26
4. Use of commands applied to validating system configuration.	27
5. Change Mappings window	30
6. Volume Properties window	31
7. Device names	42
8. Controller Properties window	44
9. Device identification information	45
10. Mappings window	46
11. Initial Automatic Discovery window	61
12. Enterprise Management window	62
13. Subsystem Management window	64
14. RVSD Cluster Configuration example with single FAST500/700 - Storage Manager v7.10 or v8.21, shown with 2 partitions (4 maximum) per FASTT	70
15. RVSD Cluster Configuration example with multiple FAST500/700 - Storage Manager v7.10 or v8.21, shown with 2 partitions per FASTT	71
16. RVSD Cluster Configuration example with multiple FAST500/700 - Storage Manager v7.10 or v8.21, shown with 1 partition per FASTT	72
17. Typical RVSD Cluster Configuration example with multiple FAST500/700 - Storage Manager v7.10 or v8.21, shown with 1 partition per FASTT	73
18. HACMP/GPFS Cluster Configuration example with single FAST500/700, shown with one partition per FASTT	74
19. HACMP/GPFS Cluster Configuration example with multiple FAST500/700, shown with two partitions per FASTT	75
20. FlashCopy repository logical drive properties	81
21. Repository capacity settings.	82
22. Listing premium features	86
23. Features list showing FlashCopy enabled.	86
24. Feature list showing FlashCopy not enabled.	87
25. FlashCopy feature icon	87
26. FlashCopy feature icon disabled	87
27. Create FlashCopy menu options	88
28. Create a FlashCopy Logical Drive Wizard startup	89
29. Wizard Introduction window	90
30. Allocate Capacity window.	91
31. Specify Logical Drive Parameters.	92
32. Preview window	93
33. Specify Array Parameters window	94
34. Specify Names window	95
35. Specify Repository Drive Capacity	96
36. Preview window	97
37. FlashCopy volumes	98
38. Undefined FlashCopy disk	99
39. Define Additional Mapping	99
40. Define Additional Mapping window	100
41. Mapped FlashCopy logical disk	101
42. FlashCopy icon states	102
43. Repository Logical Drive Properties	103
44. Base Repository Logical Drive Properties	104
45. Repository drive warning threshold.	105
46. Choosing to disable the FlashCopy drive	107
47. Disable FlashCopy confirmation window	107

48. Icon showing the disabled FlashCopy logical drive	108
49. Recreating a FlashCopy.	109
50. Deleting the FlashCopy logical drive	112
51. Delete FlashCopy Logical Drive Dialog window	112

Tables

1.	Installation sequences of Storage Manager 8.21 software packages by host type	4
2.	Supported versions of hardware for AIX systems	11
3.	Attributes for dar devices	18
4.	Attributes for dac devices.	20
5.	Attributes for hdisk devices	20
6.	Supported versions of hardware for HP-UX systems	35
7.	Primary and secondary path matrix information.	43
8.	Supported versions of hardware for Solaris systems	47
9.	JNI adapter configuration files	56
10.	Filesets required for AIX 4.3.3 RDAC	123
11.	Filesets required for AIX 5.1 RDAC	123
12.	HP-UX kernel parameter configuration requirements	125
13.	Configuration file name: /kernel/drv/fca-pci.conf	131
14.	Configuration file name: /kernel/drv/jnic146x.conf	131
15.	Configuration file name: /kernel/drv/jnic.conf	132
16.	Configuration file name: /kernel/drv/fcaw.conf	132

About this document

This document provides information about setting up, installing, configuring, and working with the IBM TotalStorage Fibre Array Storage Technology (FAST) Storage Manager Version 8.21 in UNIX[®] and AIX[™] environments.

Throughout this document, the terms *storage management software* and *Storage Manager 8.21* refer to the IBM TotalStorage FAST Storage Manager Version 8.21. Individual components of the storage management software are identified by name.

Use this document to:

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

Who should read this document

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

How this document is organized

Chapter 1, "Introduction" on page 1 provides an introduction to the Storage Manager 8.21 product, including information about product resources in addition to this document.

Chapter 2, "Installing storage management station software on AIX systems" on page 11 provides information about installing and using Storage Manager 8.21 with an AIX operating system, step-by-step instructions for installing the AIX software on a management station, and step-by-step instructions for installing the AIX management software on a host.

Chapter 3, "Installing storage management station software on HP-UX systems" on page 35 provides information about installing and using Storage Manager 8.21 with the HP-UX operating system, step-by-step instructions for installing the HP-UX software on a management station, and step-by-step instructions for installing the HP-UX management software on a host.

Chapter 4, "Installing storage management station software on Solaris systems" on page 47 provides information about installing and using Storage Manager 8.21 with the Solaris operating system, step-by-step instructions for installing the Solaris software on a management station, and step-by-step instructions for installing the Solaris management software on a host.

Chapter 5, "Completing the software installation" on page 59 provides step-by-step instructions for performing post-installation tasks.

Chapter 6, “Storage Manager 8.21 with high-availability cluster services” on page 67 provides information about the high-availability clustering system options.

Appendix A, “AIX system requirements” on page 123, Appendix B, “HP-UX system requirements” on page 125, and Appendix C, “Solaris system requirements” on page 127 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 8.21.

Appendix D, “MC Service Guard configuration details” on page 129 provides the procedure for correcting the primary and alternate paths of the imported volume groups that are changed after using **vgimport -m -s** with LVM commands.

Chapter 7, “FlashCopy and the FAStT Storage Server” on page 77 introduces FlashCopy, which is a premium feature of the IBM TotalStorage FAStT Storage Manager 8.21. It describes the various components of FlashCopy, followed by a step-by-step guide to using the features.

Appendix E, “JNI Host Bus Adapter settings” on page 131 contains JNI host bus adapter (HBA) settings.

Related publications

This section lists documents, online help systems, and Web sites. These resources provide related information that might be of interest to you.

Documents

The following documents are available in Adobe Acrobat PDF on the Storage Manager 8.21 installation CD and at the following Web site:

www.ibm.com/pc/support/

- *IBM TotalStorage FAStT Storage Manager Version 8.21 for UNIX and AIX Environments, Installation and Support Guide* (this document)
- *IBM FAStT Storage Manager Version 8.21 Installation and Support Guide for Windows NT[®] and Windows[®] 2000*
- *IBM FAStT Storage Manager Version 8.21 Installation and Support Guide for Novell NetWare*
- *IBM FAStT Storage Manager Version 8.21 Installation and Support Guide for Linux*
- *IBM TotalStorage FAStT700 Fibre Channel Storage Server Installation Guide*
- *IBM TotalStorage[™] FAStT700 Fibre Channel Storage Server User's Guide*
- *IBM FAStT700 Fibre Channel Cabling Instructions*
- *IBM FAStT500 RAID Controller Enclosure Unit Installation Guide*
- *IBM FAStT500 RAID Controller Enclosure Unit User's Reference*
- *IBM Netfinity[®] Fibre Channel Cabling Instruction*
- *IBM FAStT200 and FAStT200 HA Storage Servers Installation and User's Guide*
- *IBM FAStT200 Fibre Channel Cabling Instructions*
- *IBM TotalStorage FAStT EXP700 Storage Expansion Unit Installation and User's Guide*
- *IBM FAStT EXP500 Installation and User's Guide*
- *IBM TotalStorage FAStT C2 - 133 Host Bus Adapter Installation and User's Guide*

- *IBM FAStT FC - 2 Host Bus Adapter Installation and User's Guide*
- *IBM FAStT Host Adapter Installation and User's Guide*
- *IBM FAStT Remote Mirror Option Installation and User's Guide*
- *IBM Fibre Channel Basic SAN Configuration Setup Guide*
- *IBM FAStT MSJ User's Guide*
- *IBM Fibre Channel Problem Determination Guide*
- *IBM Fibre Channel Hardware Maintenance Manual*
- *IBM FAStT Storage Manager Script Commands*

You can also order publications through the IBM Publications Ordering System at the following Web site:

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

Additional reference

The following document is available in Adobe Acrobat PDF at the following Web site:

www.ibm.com/redbooks

IBM TotalStorage FAStT700 and Copy Services Redbook

Online help

Storage Manager 8.21 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 8.21. Click **Help** on the toolbar or press F1.

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

- For AIX, see Chapter 2, "Installing storage management station software on AIX systems" on page 11 and Appendix A, "AIX system requirements" on page 123.
- For HP-UX, see Chapter 3, "Installing storage management station software on HP-UX systems" on page 35 and Appendix B, "HP-UX system requirements" on page 125.
- For Solaris, see Chapter 4, "Installing storage management station software on Solaris systems" on page 47 and Appendix C, "Solaris system requirements" on page 127.

Web sites

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

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

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

www.ibm.com/storage/

How to send your comments

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

- E-mail

Submit your comments electronically to:

starpubs@us.ibm.com

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

- Mail or fax

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

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

Chapter 1. Introduction

The IBM TotalStorage FAStT Storage Manager Version 8.21 for UNIX and AIX environments, known as *Storage Manager 8.21*, is a set of client and host tools that allow you to manage IBM FAStT200, FAStT500, and FAStT700 storage subsystems from a storage management station.

You can install Storage Manager 8.21 on a *storage management station*, which is the system that is responsible for managing all, or a portion of, a storage network. The storage management station communicates with the network management agents that reside in the managed nodes using a network management protocol, such as Simple Network Management Protocol (SNMP). Storage management commands are sent to the storage subsystem controllers, where the controller firmware validates and runs the commands, and then returns status and configuration information to the client software.

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

This document provides system administrators with information about installing, configuring, and working with Storage Manager 8.21 in UNIX and AIX environments. Before installing Storage Manager 8.21, consult the following documentation:

readme.txt files

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

<http://www.pc.ibm.com/qtechinfo/MIGR-43793.html>

Once there, follow the link to *v8.21* information.

IBM FAStT Storage Manager Concepts Guide

Use this reference document to become familiar with the terminology and the features of the Storage Manager 8.21 software. This document is available on the installation CD and maintained at the following Web site:

<http://www.storage.ibm.com/hardsoft/disk/fastt/>

For information about installing Storage Manager 8.21 software on AIX, HP-UX, or Solaris systems, refer to this document. When you have completed the entire installation process, refer to the following online help systems, which contain information that is common to all operating system environments. You can access the help systems from the Enterprise Management and Subsystem Management windows in Storage Manager 8.21 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 8.21 features

Storage Manager 8.21 supports FAStT200, FAStT500, and FAStT700 storage servers, and provides the features that are described in this section.

Event monitoring

Event monitoring allows you to monitor storage subsystems, even when the Enterprise Management window is not open. A toolbar button and an option, **Tools** → **Update monitor**, allow updates to the monitor if changes have occurred to your storage subsystem configuration.

Heterogeneous environment

IBM TotalStorage FAStT Storage Manager Version 8.21 for UNIX and AIX environments now provides heterogeneous host support. Storage partitioning allows you a high degree of flexibility when attaching multiple hosts running different operating systems to the same storage server. Support is achieved through the use of independent and dedicated storage partitions. Heterogeneous support has the limitation of one operating system per dedicated storage partition and the constraint of 64 partitions per FAStT 700 storage server. Separate dedicated storage partitions are required because a partition will have unlimited access to all logical drives in a particular storage partition. For more detailed information about how to create and setup storage partitions please see either:

- Chapter 2, “Installing storage management station software on AIX systems” on page 11
- Chapter 3, “Installing storage management station software on HP-UX systems” on page 35
- Chapter 4, “Installing storage management station software on Solaris systems” on page 47

Configuration replication

Configuration replication allows you to save the logical configuration of a storage subsystem, and then load it on an identical storage subsystem. This feature can be used to replicate a logical configuration from one storage subsystem to another, or to save a storage subsystem configuration for backup.

Auto Volume Transfer/Auto Disk Transfer (AVT/ADT)

Auto Volume Transfer (AVT), also known as Auto Disk Transfer (ADT) and referred to in this document as AVT/ADT, is a failover method. By this method, the storage subsystem initiates the transfer of a volume, or volumes, when it detects a controller or path failure. If the host detects the failure, it automatically sends I/O to the remaining functional path, simultaneously initiating an AVT/ADT event. The alternative to this approach is the Redundant Dual Active Controller (RDAC) model. By this model, the host initiates an explicit failover for all path failure events, except a controller failure.

Command-line interface

The command-line interface is based on the script-engine commands found in the script editor and allows you to issue commands to the storage subsystems. See the SMcli.txt file on the installation CD or the command line interface (CLI) online help topic for usage information.

Controller diagnostics

You can access the controller diagnostics from the **Controller** → **Run Diagnostics** menu in the storage management software. These diagnostics allow you to test the host-side and drive-side fibre-channel loops.

Access volume mapping

Access volume mapping allows you to assign and change the LUN assignment for the access volume. An access volume is required for communicating in a host-agent-managed environment using the fibre-channel input/output (I/O) path, but is not required for a directly-managed storage subsystem that is connected through the Ethernet cable connection.

High-availability cluster services

A high-availability clustering system allows 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 local area network (LAN) component. If a component fails, the redundant component takes over cluster services and coordinates the transfer between components.

Redundant Dual Active Controller (RDAC)

The Redundant Dual Active Controller (RDAC) package, also known as the Redundant Disk Array Controller (RDAC) package, is a multipath device driver that provides controller failover support when a failure occurs anywhere along the fibre-channel I/O path.

Direct-attached configuration

Storage Manager 8.21 supports FASTT200, FASTT500, and FASTT700 storage servers in direct-attached homogenous environments.

FlashCopy

Storage Manager 8.21 supports FASTT200, FASTT500, and FASTT700 storage server use of the FlashCopy option, a premium feature of the FASTT Storage Manager 8.21. A FlashCopy logical drive is a point-in-time (PIT) image of a logical drive. It is the logical equivalent of a complete physical copy, created much more quickly than a physical copy, and requiring less disk space.

Storage Manager 8 software packages

Storage Manager 8.21 contains the following software packages:

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

Install the software packages in the sequences shown in Table 1 on page 4. For installation instructions, see the referenced sections. For an overview of each package, see the subsections that follow the table.

Table 1. Installation sequences of Storage Manager 8.21 software packages by host type

Step	AIX	HP-UX	Solaris
1	SMruntime (see “Installing SMruntime on AIX hosts” on page 13)	SMruntime (see “Installing the host software on HP-UX hosts” on page 36)	SMruntime (see “Installing the runtime software on Solaris hosts” on page 48)
2	SMclient (see “Installing SMclient on AIX hosts” on page 14)	SMclient (see “Installing the host software on HP-UX hosts” on page 36)	SMclient (on Solaris 8 only; see “Installing the client software on Solaris hosts” on page 49; supports Solaris 2.6, 2.7, and 2.8)
3	RDAC (see “Installing host RDAC software on AIX hosts” on page 14)	SMagent (see “Installing the SMagent on HP-UX hosts” on page 38)	RDAC (see “Installing host software on Solaris hosts” on page 50)
4		SMutil (see “Installing SMutil on HP-UX hosts” on page 38)	SMagent (see “Installing SMagent on Solaris hosts” on page 52)
5			SMutil (see “Installing host software on Solaris hosts” on page 50)

Storage Manager 8 client software package

The Storage Manager 8.21 client software, SMclient, is a Java™-based GUI interface that allows you to configure, manage, and troubleshoot FAST200, FAST500, or FAST700 storage servers as well as the EXP700 expansion enclosure in a storage subsystem through a host system or through a storage management station. Specifically, SMclient allows you to:

- Configure disk arrays and logical volumes
- Assign names to arrays and volume groups
- Assign logical volumes to storage partitions
- Replace and rebuild failed disk drives
- Expand the size of arrays
- Change RAID-type arrays
- Monitor the status of FAST storage servers
- Perform troubleshooting and maintenance tasks, such as downloading firmware to controllers and upgrading NVSRAM
- Expand capacity using the EXP700 1 Gb or 2 Gb fibre-channel enclosure

Additionally, SMclient with FAST700 allows you to:

- Loop four fibre-channel hosts using eight host connections with either 1 Gb or 2 Gb channels
- Assign redundant RAID controllers

The SMclient contains two main components:

Enterprise Management

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

Subsystem Management

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

For more information about this software, see either Chapter 2, “Installing storage management station software on AIX systems” on page 11, Chapter 3, “Installing storage management station software on HP-UX systems” on page 35, or Chapter 4, “Installing storage management station software on Solaris systems” on page 47 and the *IBM FASTT Storage Manager Concepts Guide*.

Storage Manager 8 agent software package

The Storage Manager 8.21 agent (SMagent) package contains the host-agent software, which you can use on HP-UX host systems to manage storage subsystems through the host fibre-channel connection. The host-agent software takes requests from a storage management station that is connected to the host through a network connection and passes the requests to the storage subsystem controllers through the fibre-channel I/O path.

For more information about managing storage subsystems through the host agent, see “Host-agent (in-band) management method” on page 6.

Storage Manager 8 utility software package

Use the Storage Manager 8.21 utility (SMutil) package to register and map new logical drives to the operating system. Install SMutil on all HP-UX and Solaris host systems attached to a storage subsystem. The host computers are attached to the storage subsystem through the fibre channel.

RDAC

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

Note: The AIX RDAC driver files are not included on the Storage Manager 8.21 installation CD; you must follow the instructions in “Installing host RDAC software on AIX hosts” on page 14 to download them from the appropriate Web site.

Storage subsystem management methods

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

- The host-agent (in-band) management method. In this method, you manage the storage subsystems through the fibre-channel I/O path to the host.

Note: You cannot use this management method on AIX systems.

- The direct (out-of-band) management method. In this method, you manage the storage subsystems directly over the network through the Ethernet connection to each controller.

Host-agent (in-band) management method

When you use the host-agent (in-band) management method, you manage the storage subsystems through the fibre-channel I/O path to the host. The management information can be processed by the host or passed to the storage management station through the network connection. Figure 1 shows the host-agent (in-band) management method.

Note: You cannot use the host-agent (in-band) management method on AIX systems.




Figure 1. Host-agent (in-band) managed storage subsystems

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

- You do not need to run Ethernet cables to the controllers.
- You do not need a Dynamic Host Configuration Protocol (DHCP) bootstrap protocol (BOOTP) server to connect the storage subsystems to the network.
- You do not need to configure the controller network (described in Chapter 2, "Installing storage management station software on AIX systems" on page 11, Chapter 3, "Installing storage management station software on HP-UX systems" on page 35, or Chapter 4, "Installing storage management station software on Solaris systems" on page 47).

- When adding devices, you need to specify a host name or Internet Protocol (IP) address for the host only, not for the individual controllers in a storage subsystem. Storage subsystems that are attached to the host are automatically discovered.


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

- You are limited to configuring one less LUN than the maximum number allowed by the operating system and host adapter that you are using.
- The host-agent requires a special logical drive, called an *access volume*, to communicate with the controllers in the storage subsystem.
- If you are upgrading controllers from firmware version 3.x to version 4.x and your host system has already configured its maximum number of LUNs, you must give up a LUN to be used as an access volume.

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 2 on page 8 shows the direct (out-of-band) management method.



SJ000708

Figure 2. Direct (out-of-band) managed storage subsystems

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

- The Ethernet connections to the controllers enable a storage management station running SMclient to manage storage subsystems that are connected to a host running one of the following operating systems:
 - AIX
 - HP-UX
 - Linux
 - Microsoft Windows NT or Windows 2000
 - Solaris
 - Any other operating systems that are supported by Storage Manager 8.21

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

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

Operating system requirements

Storage Manager 8.21 supports the FAStT200, FAStT500, and FAStT700 storage servers in the following environments:

- IBM AIX 4.3.3 or 5.1 (see Appendix A, “AIX system requirements” on page 123)
- HP-UX 11.0 (32 bit or 64 bit), HP-UX 11.11, or HP-UX 11i (see Appendix B, “HP-UX system requirements” on page 125)
- Sun Solaris 2.6, 2.7, or 2.8 (see Appendix C, “Solaris system requirements” on page 127)

Setting up IP addresses for FAStT storage controllers

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.

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 CLI commands through serial cables that are connected to a terminal emulator.

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

1. Stop all I/O to the FAStT controllers.
2. Connect a null modem serial cable from one of the controllers to a system with a terminal emulator available.

Note: Use HyperTerminal Private Edition Version 6.3 or later. Download the latest version from the following web site:

www.hilgraeve.com

3. Open the HyperTerminal and from the menu bar click **File** → **Properties** → **Configure**. Choose the following settings:
 - 115200 Baud
 - 8 Data Bits
 - 1 Stop Bit
 - No parity
 - XON/XOFF Flow Control
4. Connect to the FAStT storage server and send a break signal (Ctrl+Break for most emulators).
5. Repeat this step until the following message is displayed:

Press the space bar for baud rate within 5 seconds.

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

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

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

Note: The default IP address settings when released from manufacture are as follow:

- Controller A's IP address = 192.168.10.101
- Controller B's IP address = 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. Disconnect from the first controller and connect to the second controller.

14. Repeat steps 1 on page 9 - 13 to assign the second IP address to the second controller.

15. Restart the FAStT storage server by turning the controller unit off and on.

Chapter 2. Installing storage management station software on AIX systems

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

- Hardware and firmware requirements
 - Creating a direct-attached configuration
 - Creating a SAN-attached configuration
- AIX restrictions
- Client software installation
- Host software installation
- AIX configuration information

Hardware and firmware requirements

Table 2 lists the supported versions of hardware to use with Storage Manager 8.21.

Table 2. Supported versions of hardware for AIX systems

Product Name	Model	Product release and firmware version
IBM TotalStorage FAStT700 RAID Controller Enclosure Unit	1742-1RU	Appware 05.21.05.09, NVSRAM CNV1742R821NT012
IBM TotalStorage FAStT700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9319
IBM FAStT500 RAID Controller Enclosure Unit	3552-1RU	Appware 05.21.05.09, NVSRAM CNV3552R821NT012
IBM FAStT500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9163
IBM FAStT200 RAID Controller Enclosure Unit	3542-1RU and FC 2101	Snapware 05.20.12.09, NVSRAM CNV3542R821NT030
IBM FAStT200 RAID and Storage Unit, Double Controller	3542-2RU	Snapware 05.20.12.09, NVSRAM CNV3542R821NT030
IBM HA Emulex LP7000	FC 6227	3.22A1
IBM HA Emulex LP9000	FC 6228	3.82A1
Brocade switch	2109-S08 2109-S16	2.6.0.e
Brocade switch	2109-F16	3.0.2f
Brocade switch	2109-M12	4.0.2
Brocade switch	3534-F08	3.0.2h
McData switch	2032-064 2031-216, 2031-232 2032-140	4.01.00
INRANGE switch	2042-001 2042-128	3.2.1

Read the README file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware, NVSRAM, disk drive firmware, and host adapter device drivers:

www.pc.ibm.com/qtechinfo/MIGR-43793.html

Once there, follow the link to *v8.21* information.

Creating a direct-attached configuration

Storage Manager 8.21 supports FAStT200, FAStT500, and FAStT700 storage servers in direct-attached AIX (4.3.3 or 5.1) configurations.

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

- One or two AIX servers can be connected to the FAStT storage server.
 - FAStT200 can support one AIX server.
 - FAStT500 can support two AIX servers.
 - FAStT700 can support two AIX servers.
- Two server FAStT500 or FAStT700 configurations require four host-side mini-hubs, each with exactly one FC connection from each HBA to a mini-hub.
- There are two or four HBAs (FC 6227 or FC 6228) per FAStT storage server. Each pair must be configured to one FAStT partition.
- No external hubs are being used.

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

1. Connect the HBAs to each controller or mini-hub port of the FAStT storage server.
2. Start the system.
3. Configure and verify the configuration as shown in the next sections.

Creating a SAN-attached configuration

Storage Manager 8.21 supports FAStT200, FAStT500, and FAStT700 storage servers in a SAN environment through Fibre Channel switches in AIX configurations. 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 connected to the same switch that is connected to the same FAStT controller port.
- If connecting through a single Fibre Channel switch, such as an IBM 2109-F16, each Fibre Channel HBA and controller port must be in its own fabric zone.

Consult the documentation provided by the switch manufacturer for information about zoning for more information. Multiple FAStT devices can be configured to the same set of Fibre Channel HBAs through a FC switch.

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

1. Set the required zones on the Fibre Channel switch, if applicable.
2. Connect the HBAs to the switch or switches.
3. Connect the FAStT storage subsystems to the switch or switches.
4. Start the system.
5. Configure and verify the configuration as shown in the next sections.

AIX restrictions

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

- F-RAID Manager is not supported.
- The maximum number of partitions per AIX host, per FAStT storage server, is two.
- Each AIX host can support two or four host bus adapters (FC 6227 or 6228) and up to two FAStT storage partitions, each requiring two host bus adapter connections.
- Direct-attach configurations are restricted to single-initiator configurations only. Only one connection to each FAStT mini-hub is allowed.
- Single HBA configurations are not allowed. Each connection to a partition must have two HBAs configured.
- Single-switch configurations are allowed, but each HBA and FAStT controller combination must be in a separate SAN zone.
- All volumes that are configured for AIX must be mapped to an AIX host group. Connecting and configuring to volumes in the default host group is not allowed. See “Performing the initial configuration of storage subsystems on AIX hosts” on page 15.
- Other storage devices, such as tape devices or other disk storage, must be connected through separate HBAs and SAN zones.

Installing the host software on AIX hosts

Use the following procedure to install the host software on an AIX storage management station. Install the software in the following order:

1. SMruntime
2. SMclient
3. Update RDAC device driver

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

Prerequisites

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

- The AIX host on which you are installing the SMruntime software meets the minimum hardware and software requirements described in “Hardware and firmware requirements” on page 11 and “Software requirements” on page 123.
- The correct filesets are present on the system.

Note: If the filesets are not present, follow the instructions in “Installing host RDAC software on AIX hosts” on page 14 to download them from the appropriate Web site.

- For a list of AIX 4.3.3 filesets, see Table 10 on page 123.
- For a list of AIX 5.1 filesets, see Table 11 on page 123.

Installing SMruntime on AIX hosts

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

1. Install SMruntime by typing the following command:

```
# installp -a -d /<complete path name>/SMruntime.aix-08.20.65.02.bff SMruntime.aix.rte
```

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

```
# lspp -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 on AIX hosts

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

1. Install SMclient by typing the following command:

```
# installp -a -d /<complete path name>/SMclient.aix-08.20.65.02.bff SMclient.aix.rte
```

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

```
# lspp -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 host RDAC software on AIX hosts

After you install the Storage Manager 8.21 client software and configure your storage subsystems, use these instructions to install the appropriate Storage Manager 8.21 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 temporary fix (PTF) 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 PTFs on an AIX system and describes how to verify the current RDAC drive PTF.

Prerequisites for installing RDAC

See Table 10 on page 123 or Table 11 on page 123 for a list of PTF filesets required by version.

Verifying AIX RDAC Software

Use the **lspp** command to verify that the correct version numbers of the AIX host RDAC software is installed:

```
# ls1pp -ah filename
```

Where *filename* is one of the required filesets listed in Table 10 on page 123 or Table 11 on page 123.

For example:

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

Installing RDAC on AIX hosts

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

1. Go to one of the following Web sites:

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

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

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

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

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

```
# cfgmgr -v
```

Performing the initial configuration of storage subsystems on AIX hosts

Complete the following procedure to configure Storage Manager 8.21 for an AIX system, which can be performed from Storage Manager client running on an AIX or non-AIX system.

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

Note: See “Setting up IP addresses for FAStT storage controllers” on page 9 for information about assigning IP addresses to the controllers.

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

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

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

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

Updating FAStT firmware and NVSRAM

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

Note: Even though concurrent firmware upgrades are supported, it is recommend that I/O is quiesced before upgrading firmware. Concurrent upgrades of NVSRAM, however, are not supported.
2. Complete the following steps to upgrade the NVSRAM.
 - a. Mount and locate the NVSRAM file on the installation CD (/cdrom/NVSRAM/3542, /cdrom/NVSRAM/3552, or /cdrom/NVSRAM/1742), or download the correct version from the Web site. Place the file in a designated directory on the host system.
 - b. In the Subsystem Management window, click **Storage Subsystem** → **Download** → **NVSRAM**.
 - c. Type or select the full pathname of the NVSRAM directory.
 - d. Double-click the NVSRAM filename or click **OK** to select the correct file.
 - e. Click **Yes** to start the download.
 - f. Verify that the NVSRAM was successfully installed.
3. After you ensure that all I/O to the controllers is stopped, complete the following steps to upgrade the controller firmware.
 - a. Mount and locate the firmware file on the installation CD (/cdrom/Firmware/3542, /cdrom/Firmware/3552 or /cdrom/Firmware/1742), or download the correct version from the Web site. Place the file in a designated directory on the host system.
 - b. In the Subsystem Management window, click **Storage Subsystem** → **Download** → **Firmware**.
 - c. Type or select the full pathname of the firmware file.
 - d. Click **OK** to update both controllers.
 - e. Click **Yes** to start the download. A new window opens.
 - f. Close the Subsystem Management window and then reopen it to complete the firmware update.
 - g. Verify that the firmware was successfully installed.

Setting 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: To verify that the host port matches the AIX host, run the **lsdev -Cc adapter | grep fcs** command. Identify the fcs by associating the number of host bus adapters (HBAs), then run the **lscfg -vl fcs# | grep Network** command. Verify that the number that displays under Network Address matches the host port number in the GUI. Repeat this procedure for the second host port.

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

AIX configuration information

This section contains the following AIX configuration information:

- Viewing and setting attributes of the RDAC driver for AIX
- Identifying the controller ID numbers
- Identifying device names and bus numbers
- Identifying logical drives by operating system device names
- Redistributing volumes in case of failure
- Creating a direct-attached configuration

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

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

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

hdisk These devices represent individual LUNs on the array.

When these devices are configured, the Object Data Manager (ODM) is updated with default parameters. In most cases and for most configurations, the default parameters are satisfactory. However, there are some parameters that may be modified for maximum performance and availability. See “Recommended attribute settings” on page 24. Use the **lsattr -EI** command to view attribute settings on an AIX system, as shown in the following examples.

Definitions of attribute settings

Some attributes for dars, dacs, and hdisks, are for information purposes only, showing how the FAStT is configured or its current state. You can modify other attributes with SMIT or with the AIX **chdev -P** command. Attributes with “True” in the “Changeable” column can be modified from their default settings. “False” means that they are for informational or state purposes only. Another method of determining attributes that can be modified is with the **lsattr -EI** command. Attributes with “True” in the last column of the output can be modified. See examples below in Table 3, Table 4 on page 20, and Table 5 on page 20. Display the default with the **lsattr -DI** command.

Table 3. 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 config time by RDAC software.
<i>all_controller</i>	List of controllers that comprise this array; usually there are two dac devices.	False	Set at config time by RDAC software.
<i>held_in_reset</i>	Name of the Truecontroller 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 config time by RDAC software. (should not be changed)
<i>load_balancing</i>	Indicator that shows whether load balancing is enabled (yes) or disabled (no); see the definition of the <i>balance_freq</i> attribute for more information.	True	yes or no (see restrictions on use)
<i>autorecovery</i>	Indicator that shows whether the device returns the array to dual-active mode when it detects proper operation of both paths and controllers (yes) or not (no).	True	yes or no (see restrictions on use)
<i>lun_bitmap</i>	Bitmap that shows which controller own which LUN at the time of configuration.	True	While changeable, it is not recommended that this be changed. LUN ownership should be controlled by the Storage Manager.

Table 3. Attributes for dar devices (continued)

Attribute	Definition	Changeable (T/F)	Possible Value
<i>hlthchk_freq</i>	Number that specifies how often health checks are performed, in seconds.	True	1 - 9999 - should not be changed
<i>aen_freq</i>	Number that specifies how often polled AEN checks are performed, in seconds.	True	1 - 9999 - should not be changed
<i>balance_freq</i>	If <i>load_balancing</i> is enabled, number that specifies how often the system performs load-balancing on the array, in seconds.	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 at config time.
<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 4. 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 at config time.
<i>alt_held_reset</i>	Indicator that shows whether the alternate controller was in the held-in-reset state at the time of configuration (yes) or not (no).	False	yes or no - State of FAStT at config time.
<i>controller_SN</i>	Serial number of this controller.	False	Set by FAStT
<i>ctrl_type</i>	Type of array this controller belongs to; a value of 3542 indicates FAStT200; a value of 3552 indicates FAStT500; a value of 1742 indicates FAStT700.	False	3542, 3552, 1742 - Set by FAStT
<i>cache_size</i>	Cache size of this controller, in megabytes.	False	512, 1024 - Set by FAStT
<i>scsi_id</i>	SCSI identifier of this controller.	False	Set by SAN, reported by AIX.
<i>lun_id</i>	Logical unit number of this controller.	False	Set by FAStT
<i>utm_lun_id</i>	Logical unit number of this controller, or none if UTM (access volumes) is not enabled.	False	0 - 31. Set by FAStT Storage Manager.
<i>location</i>	User-defined location label for this controller; the system does not use this value.	True	Set by FAStT Storage Manager
<i>ww_name</i>	Fibre-channel worldwide name of this controller.	False	Set by FAStT
<i>GLM_type</i>	GLM type used for this controller.	False	high or low - Set by FAStT

Table 5. Attributes for hdisk devices

Attribute	Definition	Changeable (T/F)	Possible Value
<i>pvid</i>	AIX physical volume identifier, or none if not set.	False	Set by AIX

Table 5. Attributes for hdisk devices (continued)

Attribute	Definition	Changeable (T/F)	Possible Value
<i>q_type</i>	Queueing type for this device; must be set to simple .	False	Set by AIX, should be "simple".
<i>queue_depth</i>	Number that specifies the depth of the queue based on system configuration; reduce this number if the array is returning a BUSY status on a consistent basis.	True	1 - 64
<i>reserve_lock</i>	Indicator that shows whether the fcparray issues a SCSI Remove command every time a device is opened or when a Test Unit Ready sequence is issued by the driver (yes) or not (no).	True	yes or no
<i>write_cache</i>	Indicator that shows whether write-caching is enabled on this device (yes) or not (no); see the definition of the <i>cache_method</i> attribute for more information.	True	yes or no
<i>size</i>	Size of this volume.	False	Set by 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 5. Attributes for hdisk devices (continued)

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

Viewing ODM attributes in AIX

To view the attributes for dars, dacs, and hdisks, use the AIX `lsattr` command. The following examples demonstrate using the `lsattr -El` command, which shows the attributes that are currently set on the system. To view the default settings, use the `-Dl` flag instead of `-El`. Then, I'd remove the subsection heading "Using the `lsattr -El` command".

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

```

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

```

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

```

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

```

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

```

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

```

Changing ODM attribute settings in AIX

As discussed earlier, ODM attributes for the RDAC software and FASTT can be modified. This section shows the recommended settings for best performance and availability.

Recommended attribute settings

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

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

- hdisk:
 - cache_write
 - cache_method
 - prefetch_mult

If the changes are made through the Storage Manager client software, they will operate properly until you either reboot the host or restart cfmgmr. To make the changes permanent, you must use SMIT or the **chdev -P** command.

Attention: Controller Cache mirroring should not be disabled while `cache_write` is enabled. If this condition exists, the RDAC Software will automatically re-enable it the next time the system is reboot or when `cfmgmr` is run.

Restrictions for dar ODM attributes: The “autorecovery” attribute should only be set to **yes** in single initiator configurations. If other hosts are configured on the same set of FC links, you should not enable this function.

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 FASiT configurations with many volumes and hosts attached, this is a critical setting for high availability. Improper setting of this attribute could result in the loss of filesystems and system panics.

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

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

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

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

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

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

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

Verifying the installation and configuration of AIX hosts

After you set the (Object Data Manager) ODM attributes, FASiT volumes and configurations, and before you mount your file systems and install your applications,

use the following information to verify that all of your FASTT device names and paths are correct and that AIX recognizes your dars, dacs, and hdisks.

Initial device identification

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

```
# cfgmgr -v
```

Next, use the **lsdev -Cc disk** command to see if the RDAC software recognizes each FASTT volume as a “3542 (200) Disk Array Device”, each FASTT500 volume as a “3552 (500) Disk Array Device”, or each FASTT700 volume as a “1742 (700) Disk Array Device”. The following example shows the results of the command for a set of FASTT500 LUNs:

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

Identifying the controller ID numbers


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

To display the FASTT Controller Properties window:

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

Figure 3 on page 26 shows the Controller Properties window of the controller in slot A. This controller has a serial number of 1T04810361 and is represented by dac1.

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



SJ000709

Figure 3. 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 FASiT device, and then from each path to its associated logical drive. As mentioned earlier in this chapter, the RDAC software uses dars, dacs, and hdisks to represent the FASiT storage subsystem. dars represent the “disk array routers”; dacs represent the “disk array controllers”, and hdisks represent the “logical drives”, or FASiT “volumes”.

AIX provides the following commands to assist in determining the configuration of the FASiT storage subsystem:

lsdev

Displays devices and their characteristics. lsdev shows the state of the devices at boot time, or the last time `cfgmgr -v` was run. -

lsattr

Displays device attributes and possible values. Attributes are only updated at boot time, or the last time `cfgmgr -v` was run.

fget_config

Displays controllers and hdisks associated with a specified FASiT (dar). `fget_config` shows the current state and volume (hdisk) ownership.

There are several ways to correlate a system's configuration and monitor the state of FAStT storage subsystems. The use of commands in validating configurations is illustrated in Figure 4.




Figure 4. Use of commands applied to validating system configuration

lsdev command

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

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

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

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

The following example uses the **lsdev** command to show the status and location codes of eight FAStT200 hdisks. Notice that the location codes of the odd hdisks (1, 3, 5, 7) correspond to the same location code of dac0 in the previous example, and

the even hdisks (2, 4, 6, 8) correspond to the same location code of dac1. This means that the preferred paths for I/O for hdisk1, hdisk3, hdisk5, and hdisk7 is through dac0. The failover path would be through dac1. Conversely, the preferred path for the even-numbered hdisks would be through dac1, and failover path through dac0.

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

Another useful way of determine the full preferred path from hdisk to Fibre Channel fabric is to grep on the first two numbers of the location code, as in the following example:

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

fget_config command

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

NONE

the controller is not defined or is offline

RESET

the controller is in the reset state.

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

Using the lsattr command

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

commands.

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

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

# lsattr -El hdisk1
----- Parts removed:
scsi_id 0x11100                SCSI ID                 False
lun_id 0x000000000000000000      Logical Unit Number    False
cache_method fast_write          Write Caching method   True
```

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


Identifying logical drives by operating system device names

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

For example, if you use the **lsattr** command on the LUN named hdisk4, it provides the following information: it is a 6 GB LUN of type RAID, with a LUN ID of 2, and an IEEE volume name of 600A0B80000C1E63000000033BB39927 (see the following example). You can make a quick identification by locating the LUN ID on the far right side of the **Mappings Viewtab**.

```
# 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 0x0020000000000000      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 600A0B80000C1E63000000033BB39927 IEEE Unique volume name False
```

You can make a more exact correlation using the distinctive *ieee_volname* attribute. The value of this attribute on the AIX host is the same as the Unique Logical Drive Identifier on the FASTt storage subsystem. The **ieee_volname** and Unique Logical Drive Identifier can be found in the Change Mappings window, which is shown in Figure 5. (Click the **Mappings view** tab. Select **Logical Drive Name**. Right click and select **Change Mappings**.)




SJ000712

Figure 5. Change Mappings window

Another way to identify the system disk name based on the logical volume name is by using the Logical Drive Property window. Open this window by clicking the Storage Management window and right-clicking **Logical Drive Name** → **Properties**. Look for the **World-wide name**, **Capacity**, and **RAID level** properties, as shown in Figure 6 on page 31.

The following example shows a portion of the output of the **lsattr** command for **hdisk20**. Notice that the *ieee_volname* and *lun_id* attributes have the same values as the **World-wide name** and **Subsystem ID (SSID)** properties shown in Figure 6 on page 31, respectively.

```
# lsattr -E1 hdisk20
write_cache yes Write Caching enabled True
size 9216 Size in Mbytes False
raid_level 5 RAID Level False
scsi_id 0x210513 SCSI ID False
lun_id 0x0012000000000000 Logical Unit Number False
cache_method fast_write Write Caching method True
prefetch_mult 0 Multiple of blocks to prefetch on read True
ieee_volname 600A0B80000C1E63000000163BB39B7C IEEE Unique volume name False
```



SJ000713

Figure 6. Volume Properties window

Disk array errors

This section shows a list of possible Disk Array errors that could be reported in the AIX error log. Should an error occur, check your Storage Manager Major Event Log (MEL) to see if there is any correlation between the the host, SAN, and FASTT storage subsystem. You may also need to validate your configuration or replace defective hardware to correct the situation.

- **FCP_ARRAY_ERR1** *ARRAY OPERATION ERROR*
 - A permanent hardware error involving the disk array media.
- **FCP_ARRAY_ERR2** *ARRAY OPERATION ERROR*
 - A permanent hardware error.
- **FCP_ARRAY_ERR3** *ARRAY OPERATION ERROR*
 - A permanent error detected by the array adapter.
- **FCP_ARRAY_ERR4** *ARRAY OPERATION ERROR*
 - A temporary error within the array, communications, adapter, etc.
- **FCP_ARRAY_ERR5** *UNDETERMINED ERROR*
- **FCP_ARRAY_ERR6** *SUBSYSTEM COMPONENT FAILURE*
 - A degradation condition has occurred other than a disk drive.
- **FCP_ARRAY_ERR7** *CONTROLLER HEALTH CHECK FAILURE*
 - A health check on the passive controller has failed.
- **FCP_ARRAY_ERR8** *ARRAY CONTROLLER SWITCH*
 - A condition has occurred which caused an array controller switch to occur.

- **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 via 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 historic reasons only. FCP_ARRAY_ERR11 may be reused for a different error in the future.
- **FCP_ARRAY_ERR12** *POLLED AEN FAILURE*
 - An automatic error notification has failed.
- **FCP_ARRAY_ERR13** *ARRAY INTER-CONTROLLER COMMUNICATION FAILURE*
- **FCP_ARRAY_ERR14** *ARRAY DRIVE FAILURE*
- **FCP_ARRAY_ERR15** *CACHE BATTERY LOW/DATA LOSS POSSIBLE*
 - Note 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 will be logged in the errorlog. Do not replace the controller.
- **FCP_ARRAY_ERR16** *CACHE BATTERY CHARGE BELOW 87.5%*
 - Note 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 will be logged in the errorlog. Do not replace the controller.
- **FCP_ARRAY_ERR17** *WORLD-WIDE NAME CHANGED*
 - A controller has changed world wide names (most likely either it was replaced without placing it in reset first, or cabling was changed so that a different controller with the same SCSI id is on the loop).
- **FCP_ARRAY_ERR18** *RESERVATION CONFLICT*
 - An operation failed because the disk array volume (LUN) is reserved by another host.

Redistributing volumes in case of failure

If a failure occurs that initiates a controller failover, follow this procedure to redistribute logical drives to their preferred paths.

1. Repair or replace any faulty components.
2. Complete these steps to mark the failed controller active.
 - a. Go to the Subsystem Management window.
 - b. Right-click **controller A** or **controller B**, and then click **Change —> Mode Active**.
3. Redistribute volumes to their preferred paths by clicking **Subsystem Management —> Storage Subsystem —> Redistribute Logical Drive**.
4. Use the **fget_config** command to verify the active paths, as shown in this example:

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

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

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

- Hardware and firmware requirements
- SMruntime installation
- SMclient installation
- SMagent installation
- SMutil installation
- Storage partitioning for homogeneous environments
- Enabling multipath I/O with PV-links

Hardware and firmware requirements

For system-level hardware and software requirements, see Appendix B, “HP-UX system requirements” on page 125.

Table 6 lists the supported versions of hardware to use with Storage Manager 8.21.

Table 6. Supported versions of hardware for HP-UX systems

Product Name	Model	Product release and firmware version
IBM TotalStorage FAStT700 RAID Controller Enclosure Unit	1742-1RU	Firmware 05.21.05.09, NVSRAM CNV1742R821NT012
IBM TotalStorage FAStT700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9319
IBM FAStT500 RAID Controller Enclosure Unit	3552-1RU	Firmware 05.21.05.09, NVSRAM CNV3352R821NT012
IBM FAStT500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9163, 9165
IBM FAStT200 RAID and Storage Unit, Single Controller	3542-1RU and FC 2101	Firmware 05.20.12.09, NVSRAM CNV3542R821NT030
IBM FAStT200 RAID and Storage Unit, Double Controller	3542-2RU	Firmware 05.20.12.00, NVSRAM CNV3542R821NT030
HP FC Adapter	A5158A	B.11.00.06 (driver)
HP FC Adapter 2 Gb	A6795A	B.11.00.01 (driver)
Brocade switch	2109-S08 2109-S16	2.6.0.e
Brocade switch	2109-F16	3.0.2f
Brocade switch	2109-M12	4.0.2
Brocade switch	3534-F08	3.0.2h
McData switch	2032-064 2031-216, 2031-232 2032-140	4.01.00
INRANGE switch	2042-001 2042-128	3.2.1

If the controllers are shipped with a firmware release prior to version 05.20.08.00 and an NVSRAM release prior to version CNV3542R821NT016 (for FAStT200 storage servers), or a firmware release prior to version 05.21.01.01 and an NVSRAM release prior to version CNV3552R821NT007 (for FAStT500 storage servers), or a firmware release prior to version 05.21.01.01 and an NVSRAM release prior to version CNV1742R821NT007 (for FAStT700 storage servers), they must be upgraded. Read the readme file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware, NVSRAM, disk drive firmware, and host adapter device drivers:

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

Installing the host software on HP-UX hosts

The host software for HP-UX consists of the following three packages:

1. FAStT client (SMclient)
2. FAStT host agent (SMagent)
3. FAStT utilities (SMutil)

All three packages are dependent on SMruntime, which is a Java compiler for the SMclient, SMagent, and SMutil. The host software is dependent on Java Runtime Environment, if you do not install SMruntime you will receive errors during the software install analysis phase.

Use the following procedure 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. SMruntime
2. SMclient
3. SMagent
4. SMutil

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 125.
- Neither the SMruntime, SMclient, SMagent, or SMutil is installed on the host and you are ready to install it on this machine.

Installing SMruntime on HP-UX hosts

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

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

```
# mount -o cdcase -F cdfs /<put device pathname here> /cdrom  
# swinstall -s /cdrom/HP-UX/HPruntime_0804500
```

The Software Selection window opens and displays this message:

```
SMruntime
```

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

Note: You must exit the **sw install** 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 on HP-UX hosts

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

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

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

The Software Selection window opens and displays the following message:

```
SMclient
```

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

Note: You must exit the **sw install** 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 failures is reported, go to step 9 on page 38. 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 the client software on at least one additional storage management station or cluster server.

Installing the client software on HP-UX hosts

Use the following procedure to install the client software. Install the software in the following order:

1. SMagent
2. SMutils

Installing the SMagent on HP-UX hosts

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

1. Start the installation process:

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

The Software Selection window opens and displays this message:

```
SMagent
```

2. Highlight **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 install command before running the **swverify** command to verify installation.

8. Verify that the installation was successful:

```
# 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 the SMagent software installation on this HP-UX storage management station.

Installing SMutil on HP-UX hosts

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

1. Start the installation process:

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

The Software Selection window opens and displays this message:

```
SMutil
```

2. Highlight **SMutils**.
3. Select **Actions** → **Mark** for Installation.
4. Select **Actions** → **Install (Analysis)**.
5. Click **OK** when the analysis is finished.
6. In the conformation 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 install command before running the **swverify** command to verify installation.

8. Verify that the installation was successful:

```
# 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/swagent.log` file.

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

Configuring storage and partitioning for heterogeneous environments

After you install Storage Manager 8.21 software you need to configure the subsystem. System configuration includes configuring logical drives through the Subsystem Management window of the SMclient application.

For more information about configuring storage see *IBM TotalStorage FASTT700 and Copy Services*, an IBM Redbook publication.

See “Configuring storage subsystems” on page 60 for additional information about storage configuration.

Read the following before continuing:

- All FASTt storage servers ship with NVSRAM configured for Windows NT hosts and have an access volume set to LUN 31.
- AVT/ADT is enabled by default; leave it in that state for HP-UX hosts.
- If the controllers are shipped with a firmware release prior to version 05.20.08.00 and an NVSRAM release prior to version NV3542R810N016 (for FASTt200 storage servers), or NV3552R821NT007 (for FASTt500 storage servers), or a firmware release prior to 05.21.01.01 and an NVSRAM release prior to CNV1742R821NT007 (for FASTt700 storage servers), they must be upgraded. Read the readme file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware and the NVSRAM:

www.storage.ibm.com/fast200/supserver.htm
www.storage.ibm.com/fast500/supserver.htm
www.storage.ibm.com/fast700/supserver.htm

Note: The versions of firmware and NVSRAM that ship with this product can be later releases than those described in this document.

Configuring storage

Complete the following procedure to configure Storage Manager 8.21 for an HP-UX system.

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 8.21 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 the SMclient software 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 8.21 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**.
4. In the Subsystem Management window, click **Storage Subsystem** → **Profile** and review the summary portion of the output.

Read the readme file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware and the NVSRAM:

www.storage.ibm.com/fast200/supserver.htm
www.storage.ibm.com/fast500/supserver.htm
www.storage.ibm.com/fast700/supserver.htm

5. Complete the following steps to upgrade the NVSRAM.
 - a. Rescan the storage subsystem, and then click **Manage the Device** to return to the Subsystem Management window.
 - b. Click **Storage Subsystem** → **Download** → **NVSRAM**.
 - c. In the NVSRAM window, go to the directory where the latest NVSRAM file resides.
 - d. Type or select the full pathname of the NVSRAM file, and then click **Update** → **OK**.

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

6. Complete the following steps to upgrade the firmware.
 - a. Click **Storage Subsystem** → **Download** → **Firmware**.
 - b. Go to the directory where the firmware resides.

- c. Type or select the full pathname of the firmware file, and then click **Update** → **OK**.

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

7. Type the following command to restart the host:

```
# reboot
```

Partitioning storage

Note: Before you partition the storage, 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/td<X>
```

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

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

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

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

1. Run the hot_add command from HP-UX at the shell prompt. This command updates any new devices that are created or added. A dump is generated. When the hot_add utility 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 7. Notice that every logical drive or device and access logical unit number (LUN) is listed twice, because you have two paths to each logical drive.

```
#SMdevices
```

```
IBM FASTT Storage Manager Devices, Version 08.21.45.00
Built Wed Mar 20 00:58:59 GMT+00:00 2002
(C) Copyright International Business Machines Corporation, 2002
Licensed Material - Program Property of IBM. All rights reserved.

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

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

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

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

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

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

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

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

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

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

Figure 7. Device names

3. Determine the primary and alternate path for each logical drive. A part of the WWN of each logical drive is unique for each controller in the FASTT storage

server. If you examine the WWNs for the access volumes in Figure 7 on page 42, you will notice that they differ in only five digits, 75a54 and 75e60. The WWNs for the logical drives also differ in these five digits.

In the example in Figure 7 on page 42, the devices are viewed through the adapters c28 and c30. To determine the primary path of a specific logical drive do the following:

- a. Find the WWN for each access LUN. In this case, Access LUN 1 is associated with c28 and has the WWN of 75a54. Access LUN 2 is associated with c30 and has the WNN of 75e60.
- b. Identify the primary path by matching the logical drive WWN to an access LUN WWN. In this case, the WWN for the LUN 10 (75a54) is associated with adapter c28. Therefore, the primary path would be c28t1d2 and the alternate path would be c30t1d2.
- c. For future reference, enter this path information into a primary and secondary path matrix, similar to the one in Table 7.

Table 7. Primary and secondary path matrix information

Logical drive	Primary path	Secondary path
hp_1	c28t1d2	c30t1d2

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

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

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

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

6. Make a node by typing the following command:

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

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

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

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

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


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

```
#vgdisplay -v <vgname>
```

Identifying the controller ID numbers

Complete this procedure to identify the controller ID numbers:

1. Select a controller in the Physical View of the Subsystem Management window.
2. Click **Controller** → **Properties** → **Interfaces**. The Controller Properties window opens, as shown in Figure 8.



SJ000702

Figure 8. Controller Properties window

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

Identifying device names and bus numbers

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

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

```
# ioscan -f
```

2. The identification information for every device is displayed. See Figure 9.

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

Annotations in the table:
 - A box labeled "External bus number of the host adapter" points to the 'I' column value '2' for the **ext_bus** row.
 - A box labeled "Controller ID number" points to the 'I' column value '39' for the **target** row.
 - A box labeled "Target address" points to the 'H/W Path' value '8/12.8.0.2.0.0' for the **target** row.
 - A box labeled "LUN number" points to the 'H/W Path' value '0' for the **target** row.
 - A box labeled "Switch port ID number" points to the 'H/W Path' value '1.0' for the **disk** row with I=131.

SJ000701

Figure 9. Device identification information

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

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

<x>

Represents the controller ID number

<y>

Represents the target address

<z>

Represents the LUN number

For example, disk 130 in Figure 9 has a device name of c2t0d0.

- Determine the entries for the logical drives using the controller ID and LUN numbers. Find the **ext_bus** row above the logical drive listing and note the bus number in the **I** column.
- For each storage subsystem, use the Mappings window (see Figure 10 on page 46). Click the **Mappings View** tab in Subsystem Management window.
 - Select the storage subsystem node in the Topology View to see a list of all logical drives in the Mappings View.
 - Use the LUN numbers in the Mappings View to associate the logical drives with their device names.




Figure 10. Mappings window

Chapter 4. Installing storage management station software on Solaris systems

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

- Hardware and firmware requirements
- Runtime software installation
- Client software installation
- Host software installation
- Configuration information

Hardware and firmware requirements

Table 8 lists the supported versions of hardware to use with Storage Manager 8.21.

Table 8. Supported versions of hardware for Solaris systems

Product Name	Model	Product release and firmware version
IBM TotalStorage FAST700 RAID Controller Enclosure Unit	1742-1RU	Firmware 05.21.05.09, NVSRAM CNV1742R821NT012
IBM TotalStorage FAST700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9319
IBM FAST500 RAID Controller Enclosure Unit	3552-1RU	Firmware 05.21.05.09, NVSRAM CNV3352R821NT012
IBM FAST500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9163, 9165
IBM FAST200 RAID and Storage Unit, Single Controller	3542-1RU and FC 2101	Firmware 05.20.12.09, NVSRAM CNV3542R821NT030
IBM FAST200 RAID and Storage Unit, Double Controller	3542-2RU	Firmware 05.20.12.00, NVSRAM CNV3542R821NT030
JNI FC Adapter	FC64-1063	fcaw2.5.18.pkg
JNI FC Adapter	FCE-1063 S-bus 64 bit	JNIC.pkg, version 4.1.1.1
JNI FC Adapter	FCE-1473 S-bus 64 bit	JNIC146x.pkg, version 5.1.2
JNI FC Adapter	FCE-6410-N 33 Mhz	JNIC.pkg, version 4.1.1.1
JNI FC Adapter	FCE-6460-PCI 64 bit	JNIC146x.pkg, version 5.1.2
JNI FC Adapter	FCE2-1063 S-bus 64 bit	JNIC.pkg, version 4.1.1.1
JNI FC Adapter	FCE2-6412 66 Mhz 64 bit	JNIC.pkg, version 4.1.1.1
JNI FC Adapter	FCE-1063	fca-pci.2.5.18.pkg
Brocade switch	2109-S08 2109-S16	2.6.0.e
Brocade switch	2109-F16	3.0.2f

Table 8. Supported versions of hardware for Solaris systems (continued)

Product Name	Model	Product release and firmware version
Brocade switch	2109-M16	4.0.0.a
Brocade switch	3534-F08	3.0.2h
McData switch	2032-064 2031-216, 2031-232 2032-140	4.01.00
INRANGE switch	2042-001 2042-128	3.2.1

Read the readme that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware, NVSRAM, disk drive firmware, and host adapter device drivers:

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

For system-level hardware and software requirements, see Appendix C, “Solaris system requirements” on page 127.

Installing the runtime software on Solaris hosts

Use the following procedure to install the runtime software (SMruntime) on a Solaris storage management station.

Prerequisites

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

- This is the SUN SPARCstation workstation that you have identified as the storage management station.
- This machine is running Solaris 2.6, 2.7, or 8 and meets the minimum hardware and software requirements described in Appendix C, “Solaris system requirements” on page 127.
- SMruntime is not installed and you are ready to install it on this machine.

Procedure

In the following procedure, the installation CD is mounted at /cdrom/sm821. Adjust these instructions as required for your specific installation location. No restart is required during the client installation process.

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

```
pkgadd -d /cdrom/SM821/SUN_SParc/SM8runtime-Sparc<filename>.pkg
```

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

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

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

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

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

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

Installing the client software on Solaris hosts

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

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

Prerequisites

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

- This is the Sun SPARCstation workstation that you have identified as the storage management station.
- This machine is running Solaris 2.6, 2.7 or 8 and meets the other minimum hardware and software requirements described in Appendix C, “Solaris system requirements” on page 127.
- SMruntime is installed.
- SMclient is not installed and you are ready to install it on this machine.

Installing SMclient on Solaris hosts

In the following procedure, the installation CD is mounted at /cdrom/sm821. Adjust these instructions as required for your specific installation location. No restart is required during the client installation process.

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

```
# pkgadd -d /cdrom/SM821/SUN_Sparc/SM8client-Sparc<filename>.pkg
```

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

```
The following packages are available:

1 SMclient                IBM FASTT Storage Manager 8 Client
                          (sparc) <version number>

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

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

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

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

[y, n, ?]
```

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

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

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

```
# pkgchk SMclient
```

- The installation was successful if no output is displayed on the screen.
6. You are finished with the client software installation on this Solaris storage management station. Remove the installation CD from the CD-ROM drive.
 7. Choose one of the following options:

Storage management software is required on additional storage management stations

Repeat step 1 on page 49 through step 6, using the corresponding installation profile for each storage management station as a guide.

Storage management software installation is completed on all storage management stations

You are now ready to install the host software. Go to “Installing host software on Solaris hosts”.

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.

Installing host software on Solaris hosts

After you install the Storage Manager 8.21 client software and configure your storage subsystems, use these instructions to install the appropriate Storage Manager 8.21 host software.

All Solaris hosts in your storage subsystem must have the following host software packages installed on them:

- RDAC

- SMagent
- SMutil

This section contains the procedure for installing the software.

Prerequisites

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

- This is the Sun SPARCstation workstation that you have identified as the storage management station.
- This machine is running Solaris 2.6, 7, or 8, and meets the minimum hardware and software requirements described in Appendix C, “Solaris system requirements” on page 127.
- The RDAC, SMagent, and SMutil software packages are not installed and you are ready to install them on this machine.

Note: If you plan to run Storage Manager 8.21 in an environment that includes storage arrays that you are not upgrading to 8.21, you must upgrade your existing storage manager to at least version 6.22. The Storage Manager 8.21 RDAC package replaces the version 6.22 RDAC files and is shared by version 6.22 and version 8.21 of the storage management software.

See the Storage Manager 6.x installation instructions for information about uninstalling the software.

Installing RDAC on Solaris hosts

In the following procedure, the installation CD is mounted at /cdrom/SM821. Adjust these instructions as required for your specific installation location.

1. Insert the Solaris installation CD in the CD-ROM drive.
2. Type the following command to start installing the RDAC package:

```
# pkgadd -d /cdrom/SM821/SUN_Sparc/SM8rdac-Sparc-<filename>.pkg
```

The installation process begins.

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

```
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.
```

Note: Make sure 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 “Variable settings for JNI adapter cards” on page 56.

Installing SMagent on Solaris hosts

Perform the following procedure to install SMagent:

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

```
pkgadd -d /cdrom/SM821/SUN_SParc/SM8agent-Sparc-<filename>.pkg
```

The installation process begins.

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

```
The following packages are available:

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

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

2. Type the value of the package that you are installing and press Enter.

The installation process begins. The following prompt is displayed.

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

3. Type `y` and press Enter.

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

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

4. Move to the root directory:

```
# cd /
```

5. Stop the system:

```
# halt
```

6. Restart the system with either this command:

```
# boot -r
```

or this command:

```
# boot <bootdevices> -r
```

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

You have finished installing the Storage Manager 8.21 software on this Solaris host.

Installing SMutil on Solaris Host

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

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

The installation process begins.

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

```
The following packages are available:  
  
SMutil IBM FAST 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.

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

Performing the initial configuration of storage subsystems on Solaris hosts

Read the following before continuing:

- All FASTt storage servers ship with NVSRAM configured for Windows NT hosts and have an access volume set to LUN 31. This should not present any problems for Solaris in recognizing the storage subsystem.
- If the controllers are shipped with a firmware release prior to version 05.20.08.00 and an NVSRAM release prior to version CNV3542R821NT016 (for FASTt200 storage servers), or a firmware release prior to version 05.21.01.01 and an NVSRAM release prior to version CNV3552R821NT007 (for FASTt500 storage servers), or a firmware release prior to version 05.21.01.01 and an NVSRAM release prior to version CNV1742R821NT007 (for FASTt700 storage servers), they must be upgraded.

Read the README file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware and the NVSRAM:

www.storage.ibm.com/fast200/supserver.htm

www.storage.ibm.com/fast500/supserver.htm

www.storage.ibm.com/fast700/supserver.htm

Note: Versions of firmware and NVSRAM that ship with this product can be later releases than those described in this document.

If you do not have the latest versions, download them and continue with step 1. If the version numbers are current, then go to step 6.

Complete the following procedure to configure Storage Manager 8.21 for a Solaris system:

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

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

6. Rescan the storage subsystem, and then click **Manage the Device** to return to the Subsystem Management window.
7. In the Subsystem Management window, click the **Mappings View** tab.
8. In the Mappings window, create a new host group by clicking **Mappings** → **Define** → **Host Group**.

9. Type the name of the new host group (for example, Solaris). Click **Add**, and then click **Close**.
10. Select the new host group and click **Mappings** → **Define** → **Host**.
11. 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**.
 - d. Choose the host port for the second HBA and click **Add**, and then click **Close**.
 - e. To view the added LUNs, click the host group, and then click **Mappings** → **Define** → **Define Additional Mapping**.
 - f. Type the following command to run the hot_add utility:

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

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

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

```
# format
```

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

Uninstalling Storage Manager 8.21 software

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

Default partitioning for Solaris devices

After you install the Storage Manager 8.21 software and configure the subsystems and logical drives through the Subsystem Management window of the SMclient application, run the `/etc/raid/bin/hot_add` utility from a shell prompt. (For additional information about the hot_add utility see the applicable man page.)

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

```

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

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

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

Configuring JNI host bus adapter cards

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

- For FCI-1063 cards, `/etc/raid/bin/genfcaconf`
- For all `jnic.conf` and `jnic.146x.conf` cards, `/etc/raid/bin/genjnicconf`
- For FC64-1063 and all other cards, `/etc/raid/bin/genscsiconf`, which calls other scripts as appropriate

For additional information, see the man pages for each script.

Variable settings for JNI adapter cards

Table 9 lists the configuration file for each JNI adapter driver package.

Table 9. JNI adapter configuration files

JNI HBA model numbers	JNI adapter driver package	Configuration file
FC64-1063-N	<code>fcaw2.5.18.pkg</code>	<code>/kernel/drv/fcaw.conf</code>
FCI-1063-N	<code>fca-pci.2.5.18.pkg</code>	<code>/kernel/drv/fca-pci.conf</code>
FCE2-1063, FCE2-1063, FCE-6410, FCE2-6412	JNIC.pkg, version 4.1.1.1	<code>/kernel/drv/jnic.conf</code>
FCE-1473, FCE-6412 (2 Gb)	JNIC146x.pkg, version 5.1.2	<code>/kernel/drv/jnic146x.conf</code>

Note: See information located in Appendix E, “JNI Host Bus Adapter settings” on page 131 to verify your HBA settings.

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

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

Note:

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

In all configuration files, set the following variables to the specified values:

- *target0_hba*, set to *jnic0* in the *jnic.conf* file, or to *jnic146x0* in the *jnic146x.conf* file.
- *target_wwpn*, set to the *wwpn* of the first controller (see the */var/adm/messages* file).
- *target1_hba*, set to *jnic1* in the *jnic.conf* file, or to *jnic146x1* in the *jnic146x.conf* file, *fca-pci0* in the *fca-pci.conf* file and *fcaw0* in the *fcaw.conf* file, *fca-pci1* in the *fca-pci.conf* file and *fcaw1* in the *fcaw.conf* file.
- *target1_wwpn*, set to the *wwpn* of the second controller (see the */var/adm/messages* file).

If you made any changes to the configuration file, type the following command to halt the host system:

```
# halt
```

When the OK prompt displays, type the following command to restart the host system:

```
# boot -r
```

Setting the Solaris system host type

After you install the Storage Manager 8.21 software, you must set the host type to Solaris for the operating system to work properly with the storage subsystem. This should have been done automatically, but you should verify that it occurred.

Click on the **Mappings View** tab and select **Mappings** → **Show all Host Port Information**.

In the Mappings window:

1. Create a new host group.
2. Add the host to the newly created host group.
3. Add each port to the host and ensure that the host type is set to Solaris (Sparc).

Chapter 5. Completing the software installation

This chapter describes how to complete the installation of Storage Manager 8.21 by performing these tasks:

- Failover protection
- Changing the host type
- Understanding the implications of using heterogeneous hosts in your storage subsystem
- Configuring the storage subsystem
- Starting the Subsystem Management window

Failover protection

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

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

Changing the host type

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

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

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

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

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

When you have selected the host type, continue with “Configuring storage subsystems” on page 60.

Heterogeneous hosts overview

The heterogeneous hosts feature enables hosts that are running different operating systems to access a single storage subsystem. In previous releases of Storage Manager, only hosts running the same operating system could access a single storage subsystem. Storage Manager 8.21 supports up to 64 storage partitions, which enables a multiple host-type subsystem to share storage capacity, consolidate storage, and reduce storage management costs.

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

In a heterogeneous environment, you must set each host type to the appropriate operating system during host-port definition so that the firmware on each controller can respond correctly for the operating system for that host. Before you begin setting up your heterogeneous host's configuration, see the *IBM FAS*St*T Storage Manager Concepts Guide* available on your installation CD or at the following Web site:

www.ibm.com/pc/support/

Configuring storage subsystems

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

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


Starting the Enterprise Management window

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

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

A splash screen is displayed while the client software starts. When the client software has been loaded, the Enterprise Management window and the Initial Automatic Discovery window opens (see Figure 11 on page 61).

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



SJ000704


Figure 11. Initial Automatic Discovery window

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

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

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

When the initial automatic discovery is finished, all attached hosts and attached storage subsystems are displayed in the Enterprise Management window (see Figure 12 on page 62).



SJ000705

Figure 12. Enterprise Management window

5. Make sure that all the attached hosts and storage subsystems are displayed as expected. If not, do the following:
 - a. Check the hardware and connections for possible problems (see the hardware documentation for specific procedures).
 - b. See the Enterprise Management window help topic on discovering storage subsystems and take the appropriate action that is indicated.
 - c. Make sure 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.
 - d. If a storage subsystem is duplicated in the device tree after an automatic discovery, remove the duplicate storage subsystem icon from the device tree by clicking **Edit** → **Remove Device** in the Enterprise Management window.
6. Make sure that the status of each storage subsystem is **Optimal**.
 - If any device shows a status of **Unresponsive**, remove the device from the management domain, and then add it again. See the Enterprise Management window help topic on removing and adding devices.
 - If the device still shows an unresponsive status, contact your customer service representative.
7. Go to “Configuring alert notifications”.

Configuring alert notifications

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

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

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

www.ibm.com/pc/support/

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

1. Make sure that the installation CD is inserted in the CD-ROM drive on your designated NMS.
2. From the installation CD, copy the SM8_21.MIB file from the SM8mib directory to the NMS.
3. Follow the steps required by your NMS to compile the MIB. For details, contact your network administrator or see the documentation for the storage management product you are using.
4. Go to “Starting the Subsystem Management window”.

Starting the Subsystem Management window

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

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

The Subsystem Management window for the selected storage subsystem opens (see Figure 13 on page 64).

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




Figure 13. Subsystem Management window

Renaming storage subsystems

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

Performing optional storage subsystem management tasks

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

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

If you are installing the storage management software in a cluster environment, repeat all applicable client software installation procedures for a second storage management station and for all applicable host software on a second host. See Chapter 2, “Installing storage management station software on AIX systems” on page 11, Chapter 3, “Installing storage management station software on HP-UX

systems” on page 35, or Chapter 4, “Installing storage management station software on Solaris systems” on page 47 for installation instructions.

Chapter 6. Storage Manager 8.21 with high-availability cluster services

The high-availability clustering services provided by Storage Manager 8.21 allow application services to continue when a hardware or software failure occurs. This system protects you from software failures as well as from the failure of a CPU, disk, or LAN component. If a component fails, its redundant partner component takes over cluster services and coordinates the transfer between components.

General information

Storage Manager 8.21 is certified for use with the following cluster services:

- MC/Service Guard versions A.11.09 and A.11.13 on HP-UX systems
- Veritas Cluster Server 1.3 and 2.0 on Solaris systems
- Veritas Volume Manager 3.2 on Solaris systems

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

Prerequisites for HP-UX

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

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

Prerequisites for Solaris and Veritas Cluster Server

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

General hardware requirements

Each Solaris system in the cluster requires the following hardware:

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

System dependencies

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

Veritas Volume Manager and Data Multipathing

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

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

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

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

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

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

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

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

RDAC IDs

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

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

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

2. Modify the `Rdac_HotAddIDs` line as follows:

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

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

Single points of failure

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

Storage Manager 8.21 with General Parallel File System (GPFS)

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

Minimum AIX, PSSP, HACMP, and GPFS software requirements

- PSSP 3.4, APAR IY31786
- GPFS 1.5

FASTT software and firmware requirements

- FASTT Storage Manager v7.10 or v8.21

Configuration limitations

- No direct connection allowed between the host node and a FASTT. Only switched fabric connection is allowed.
- Each server must have two or four Fibre Channel connections to each attached FASTT; one connection per FASTT controller host-side minihub; two connections per FASTT partition with each connection on different controllers. These connections must be made using two or four Fibre Channel host adapters. Multiple FASTT storage servers can be attached utilizing these same host adapters (fan out) or by adding additional HBA pairs for each partition.
- A maximum of 4 partitions are allowed per FASTT.
- RVSD Clusters can support a maximum of 2 IBM Virtual Shared Disk and RVSD servers per FASTT partition.
- HACMP/GPFS clusters can support 2 - 32 servers per FASTT partition. In this environment be sure to read and understand the AIX device drivers queue depth settings as described in Chapter 2, “Installing storage management station software on AIX systems” on page 11.
- For maximum availability, distribute the HBA and FASTT connections across separate FC switches to minimize the effects of a SAN fabric failure.
- Single Node Quorum is not supported in a dual-node GPFS cluster with FASTT disks in the configuration.

Other usage notes specific to PSSP/GPFS environments


In GPFS file systems the following FASTT cache settings are supported:

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

The performance benefits of read or write caching is application dependent.

GPFS/PSSP logical configuration diagrams


The diagrams that follow show both preferred and fail-over 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.



Note:
 The preferred path to a volume is determined by controller assignment of Volume.
 Volumes can be spread across both controllers (A & B), allowing for static load-balancing.

SJ000886

Figure 14. RVSD Cluster Configuration example with single FAST500/700 - Storage Manager v7.10 or v8.21, shown with 2 partitions (4 maximum) per FAST




Note:

The preferred path to a volume is determined by controller assignment of Volume.

Volumes can be spread across both controllers (A & B), allowing for static load-balancing.

Please note zoning requirements when connecting multiple HBAs to the same FASTt from the same server. SJ000887

Figure 15. RVSD Cluster Configuration example with multiple FAST500/700 - Storage Manager v7.10 or v8.21, shown with 2 partitions per FASTt




Note:

The preferred path to a volume is determined by controller assignment of Volume.

Volumes can be spread across both controllers (A & B), allowing for static load-balancing.

SJ000888

Figure 16. RVSD Cluster Configuration example with multiple FAST500/700 - Storage Manager v7.10 or v8.21, shown with 1 partition per FAST



Note:

The preferred path to a volume is determined by controller assignment of Volume.

Volumes can be spread across both controllers (A & B), allowing for static load-balancing.

SJ000889

Figure 17. Typical RVSD Cluster Configuration example with multiple FASt500/700 - Storage Manager v7.10 or v8.21, shown with 1 partition per FASt




Figure 18. HACMP/GPFS Cluster Configuration example with single FASt500/700, shown with one partition per FASt




Figure 19. HACMP/GPFS Cluster Configuration example with multiple FAST500/700, shown with two partitions per FAST

Chapter 7. FlashCopy and the FASSt Storage Server

This chapter introduces FlashCopy, a premium feature of the FASSt Storage Manager 8.21 and describes its various components, what they do, and how to set them up. This is followed by a step-by-step guide to using several FlashCopy features. Refer to the *IBM TotalStorage FASSt700 and Copy Services Redbook* for information about using automated scripts, the command line interface, examples and where they can be useful. You will find this redbook at:

<http://www.ibm.com/redbooks>

The FlashCopy option is available for purchase separately from IBM or an IBM Business Partner. FlashCopy feature is available only with IBM FASSt Storage Manager Version 8.21 and above.

Overview of FlashCopy

A FlashCopy logical drive is a point-in-time (PIT) image of a logical drive. It is the logical equivalent of a complete physical copy, but you create it much more quickly than a physical copy; it also requires less disk space. In FASSt Storage Manager 8.21, the logical drive from which you are basing the FlashCopy, called the *base logical drive*, must be a standard logical drive in your storage subsystem. FlashCopy can enable an application, for example a backup application, to access the FlashCopy and read the data while its base logical drive remains online and user-accessible. When the backup completes, the FlashCopy logical drive is no longer needed.

It is also possible to create several flash copies of a base logical drive and write data to the FlashCopy logical drives to perform testing and analysis. Before upgrading a database management system, for example, you might use FlashCopy logical drives to test different configurations. Then you can use the performance data that is provided by the storage management software to help you decide how best to configure your live database system.

When you make a FlashCopy, the controller suspends I/O to the base logical drive for a few seconds. Meanwhile, it creates a physical logical drive called the FlashCopy *repository logical drive* to store FlashCopy meta data and copy-on-write data. When the controller finishes creating the FlashCopy repository logical drive, I/O write requests to the base logical drive can continue. However, before a data block on the base logical drive is modified, a copy-on-write occurs. This copies the contents of blocks that are to be modified into the FlashCopy repository logical drive for safekeeping. Because the FlashCopy repository logical drive stores copies of the original data in those data blocks, further changes to those data blocks write directly to the base logical drive without another copy-on-write. And, because the only data blocks that are physically stored in the FlashCopy repository logical drive are those that have changed since the time of the FlashCopy, the FlashCopy technology uses less disk space than a full physical copy.

When you create a FlashCopy logical drive, you specify where to create the FlashCopy repository logical drive, its capacity, and other parameters. You can disable the FlashCopy when you are finished with it, for example, after a backup completes. Then, you can recreate the FlashCopy the next time you do a backup and reuse the same FlashCopy repository logical drive. Using the Disable FlashCopy and Recreate FlashCopy menu options provides a shortcut to create a new FlashCopy logical drive of a particular base logical drive. You do not need to

create a new FlashCopy repository logical drive. You can also delete a FlashCopy logical drive, which also deletes an associated FlashCopy repository logical drive.

The storage management software provides a warning message when a FlashCopy repository logical drive nears a user-specified threshold (a percentage of its full capacity, the default is 50%). When this threshold is met, storage management software allows you to expand the capacity of a FlashCopy repository logical drive using free capacity on the array. If no free capacity is available on the array, you can optionally add unconfigured capacity to the array with which to expand the FlashCopy repository logical drive.

Enabling FlashCopy

You enable the IBM FAStT Storage Manager Version 8.21 FlashCopy feature using its accompanying Key CD. This CD contains a GenKey program that you must run to enable the FlashCopy feature.

Note: The GenKey program is licensed for use on only one storage subsystem. Purchase a FAStT FlashCopy Premium Feature Key CD for each storage subsystem on which you want to enable the FAStT FlashCopy feature.

To enable FlashCopy you must perform the following tasks.

1. Obtain the feature enable identifier.
2. Generate the feature key file.
3. Use the feature key file to enable the premium feature.

Obtaining the feature enable identifier

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

1. Click **Start** → **Programs** → **FAStT Storage Manager 8.21 Client**. The Enterprise Management window opens.
2. In the left-hand pane, double-click the storage subsystem for which you want to enable the FlashCopy 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

Complete the following steps to generate the feature key file.

1. Insert the IBM FAStT Storage Manager Version 8.21 FlashCopy Premium Feature Key CD into the CD-ROM drive.
2. From a command prompt, change to the root directory on the CD-ROM drive; then, type:

```
GenKey <full pathname and filename>  
<feature enable identifier>
```

where:

- *full pathname* is the location at which you want to create the feature key file
- *filename* is the feature key file name
- *feature enable identifier* is the feature enable identifier you recorded in “Obtaining the feature enable identifier” on page 78

For example, if you want to create a feature key file named “flashcopy.key” on drive C, using the feature enable identifier “67890875432145678976543212345678”, type:

```
GenKey c:\flashcopy.key 67890875432145678976543212345678
```

Note: Always specify a feature key file name with the .key extension. The premium feature user interface uses this extension to filter the files in a directory.

3. Press Enter. A feature key file is created in the location that you specified in step 2 on page 78.

Using the feature key file to enable FlashCopy

Complete the following steps to enable FlashCopy.

1. From the Subsystem Management window, click **Storage Subsystem** → **Premium Features** → **Enable**. The Select Feature Key File window opens.

Note: The Select Feature Key File program filters files from the .key extension.

2. Select the folder in which you placed the generated key file.
3. Select the appropriate key file.
4. Click **OK**.
5. The Enable Feature window opens. Click **Yes**.
6. The FlashCopy Option is now enabled. The icon in the Premium Feature status area no longer displays a red slash. To further verify the status of the option, click **Storage Subsystem** → **Features** → **List**

Your system is now prepared to schedule use of the FlashCopy feature.

FlashCopy parameters

When you create a FlashCopy logical drive, you specify where to create the FlashCopy repository logical drive, its capacity, threshold level warning, and other parameters. The FlashCopy repository logical drive capacity is created as a percentage of the base logical drive and contains the copy-on-write data. The Storage Management software provides a warning message when your FlashCopy repository logical drive exceeds the threshold level.

The FlashCopy repository logical drive’s failure policy determines what happens when the FlashCopy repository logical drive becomes full (that is, all of its capacity has been used). The failure policy can be set to either fail the FlashCopy (default setting) or fail incoming I/O to the FlashCopy’s base logical drive.

Do not ignore the FlashCopy repository logical drive “threshold exceeded” notification. This is the last and only warning you will receive before the FlashCopy repository logical drive actually becomes full. You have the options either to increase the capacity of the FlashCopy repository logical drive or increase the FlashCopy repository logical drive threshold capacity warning level. By increasing

the warning threshold, you reduce the amount of time you will have to respond in the event of a reaching a threshold and receiving a "threshold exceeded" notification.

If a FlashCopy logical drive or FlashCopy repository logical drive is displayed as a missing logical drive, the storage subsystem has either detected drives associated with the FlashCopy, or FlashCopy repository logical drives that are no longer accessible. Missing logical drives, in most cases, are recoverable. See "Missing logical drives" on page 120.

FlashCopy logical drives


The default FlashCopy repository logical drive capacity is initially set to 20% of the base logical drive, if enough free capacity exists to create a FlashCopy repository logical drive of this size. The default threshold level for the FlashCopy repository logical drive is initially set to 50%.

If you are not sure how large to make the FlashCopy repository logical drive or how high to set the *FlashCopy repository logical drive full* warning, accept the default settings. You can estimate later how quickly the repository capacity is being used. For more information, see "Estimating FlashCopy repository life" on page 84.

Viewing FlashCopy logical drive failure settings

Complete the following steps to see the current FlashCopy logical drive failure settings.


1. Select a FlashCopy repository logical drive in the Logical view of the Subsystem Management window.
2. Click **Logical Drive** → **Properties**, or right-click and select **Properties**. The FlashCopy Repository Logical Drive - Properties window opens. See Figure 20 on page 81..



SJ000825

Figure 20. FlashCopy repository logical drive properties

3. Click the **Capacity** tab to view the currently defined settings. See Figure 21 on page 82.



SJ000826

Figure 21. Repository capacity settings

If the FlashCopy repository logical drive is set to fail, the data is not recoverable when the drive becomes full and FlashCopy cannot be accessed. The only available option (if this situation occurs) is to delete the FlashCopy logical drive or re-create the FlashCopy logical drive to create a new point-in-time image.

If the FlashCopy repository logical drive is set to fail when writing to the base logical drive, the data is recoverable. But, the FlashCopy repository logical drive capacity must be increased before writes to the base logical drive are not rejected. See "Resizing a FlashCopy repository drive" on page 109.

Note the following points:

- Deleting a FlashCopy logical drive automatically deletes the associated FlashCopy repository logical drive.
- Deleting a FlashCopy repository logical drive automatically deletes the associated FlashCopy logical drive.
- Deleting a FlashCopy logical drive and then creating it again forces you to stop the host application and unmount the base logical drive while the FlashCopy is being created again.
- Recreating a FlashCopy logical drive alleviates the need to create a FlashCopy repository logical drive, as well as remap the assigned logical drive-to-LUN mappings between the FlashCopy logical drive and the host.
- After you recreate the FlashCopy logical drive, you can change parameters on the FlashCopy repository logical drive through the appropriate menu options.

- To avoid another *FlashCopy repository logical drive full* failure, increase the capacity of the FlashCopy repository logical drive. See “Resizing a FlashCopy repository drive” on page 109.

Estimating FlashCopy repository logical drive capacity

When a FlashCopy logical drive is created, a physical logical drive called the *FlashCopy repository logical drive* is also created to store FlashCopy data and copy-on-write data. The default setting for the FlashCopy repository logical drive capacity is 20% of the base logical drive’s capacity. In general, this capacity should be sufficient. Use the following information to help determine the appropriate capacity of the FlashCopy repository logical drive:

- A FlashCopy repository logical drive can be no smaller than 8 MB.
- The amount of write activity to the base logical drive after the FlashCopy logical drive has been created dictates the size of a FlashCopy repository logical drive. As the amount of write activity to the base logical drive increases, the number of original data blocks that need to be copied from the base logical drive to the FlashCopy repository logical drive also increases.
- The estimated life expectancy of the FlashCopy logical drive contributes to determining the capacity of the FlashCopy repository logical drive. The longer a FlashCopy logical drive remains enabled, the FlashCopy repository logical drive runs the increased risk of reaching its maximum capacity. For more information, see “Estimating FlashCopy repository life” on page 84.
- The amount of management overhead required on the FlashCopy repository logical drive to store FlashCopy logical drive data contributes to determining the capacity of the FlashCopy repository logical drive. The amount of management overhead actually required is fairly small and can be calculated using the simple formulas detailed in the following section.
- There is not necessarily a one-to-one correlation between the number of data blocks that change on the base logical drive and the amount of copy-on-write data that is stored on the FlashCopy repository logical drive. Depending on the location of data blocks that need to be copied, for performance reasons, the controller might copy over a full set of 32 blocks, even if only one set of blocks has changed. Keep this in mind when determining the percentage of the capacity of the base logical drive that can be copied to the FlashCopy repository logical drive.

Calculating expected overhead

Use the following formula to calculate the amount of management overhead that is required to store FlashCopy data on the FlashCopy repository logical drive. This formula should be used merely as a guide, and FlashCopy repository logical drive capacity should be re-estimated periodically.

Note: Conversion from bytes to kilobytes, and then to megabytes, is required for this formula.

The formula to calculate the amount of management overhead required is:

$$192 \text{ KB} + (X / 2000)$$

Where *X* is the capacity of the base logical drive in bytes.

Example

For a 5 GB base logical drive, where 30% of the data blocks are expected to change on the base logical drive, the estimated FlashCopy repository logical drive capacity can be calculated as follows:

1. Convert the capacity of the base logical drive to bytes.
When converted, 5 GB equals 5 368 709 120 bytes.
2. Divide the capacity of the base logical drive (in bytes) by 2000.
When divided, the result is 2,684,354.56 bytes.
3. Convert the result from step 2 (in bytes) to kilobytes (KB).
When converted, the result is 2621.44 KB.
4. Add 192 KB to the results from step 3.
 $192 \text{ KB} + 2621.44 \text{ KB} = 2813.44 \text{ KB}$
5. Convert the result from step 4 to megabytes (MB).
When converted, the amount of management overhead required is 2.75 MB (or 0.002686 GB).
6. In this example, 30% of the data blocks on the base logical drive are expected to change. To accurately calculate the FlashCopy repository logical drive capacity, sufficient space needs to be allocated for the copy-on-write data as well as the management overhead (calculated in step 5).
To calculate the copy-on-write space required, calculate the percentage of the base logical drive expected change:
 $30\% \text{ of } 5 \text{ GB} = 1.5 \text{ GB}$
The final estimated FlashCopy repository logical drive capacity for this example is:
 $1.5 \text{ GB} + 0.002686 \text{ GB} = 1.502686 \text{ GB}$
7. In the Create FlashCopy Logical Drive Wizard: Specify Repository Capacity window, use the percent (%) full box field to set the estimated FlashCopy repository logical drive capacity. See Figure 21 on page 82.

Note: The percentage (%) full box sets the FlashCopy repository logical drive capacity as a percentage of the base logical drive. Using the percent (%) full box field, increase or decrease the percentage until the FlashCopy Repository Logical Drive Capacity value matches the estimated capacity calculated in step 6. (Some rounding up might be required.)

Estimating FlashCopy repository life

When you create a FlashCopy logical drive, you are asked to define various properties for the FlashCopy repository logical drive, including the name of the FlashCopy repository logical drive, its capacity, a logical drive-to-LUN mapping, and the repository full condition. When defining the properties of the FlashCopy repository logical drive, keep in mind the kind of usage you have planned for the FlashCopy logical drive. Understanding how the FlashCopy logical drive will be used can help you to estimate the life expectancy of the FlashCopy repository logical drive.

If numerous I/O requests are written to the base logical drive, the FlashCopy repository logical drive (containing the FlashCopy data and copy-on-write data) could eventually exceed the base logical drive capacity if all the original data blocks are changed. The default setting suggests 20% of the base logical drive capacity, but this is a setting that can be adjusted after some usage data becomes available.

The following procedure describes how to estimate the life expectancy of a FlashCopy repository logical drive. Use it merely as a guide. Note that the life expectancy of a FlashCopy repository logical drive should be carefully re-estimated periodically.

1. Select the FlashCopy repository logical drive in the Logical view. Click **Logical Drive** → **Properties** and click the **Capacity** tab. Alternately, you can right-click the FlashCopy repository logical drive, select **Properties**, and click the **Capacity** tab.
2. Select the FlashCopy repository logical drive in the Logical view. Click **View** → **GoTo** → **FlashCopy Logical Drive**. Alternately, right-click the FlashCopy repository logical drive and select **Go To FlashCopy Logical Drive**.
3. Record the creation timestamp day and time.
4. Record the capacity used (GB) and the available capacity (GB).
5. Determine the elapsed time (t) by subtracting the creation time from the current time, expressing the elapsed time in either minutes, hours, or days.
6. The total time (Tr) that the repository is available for copy-on-write data can now be estimated (based on the current usage) by multiplying the elapsed time (t) by the available capacity (Ct), and then dividing the resultant number by the capacity used (Cu).

The formula you need to use is:

$$Tr = (t * Ct) / Cu$$

Where:

Tr	Total time available
t	Elapsed time
Ct	Available capacity
Cu	Capacity used

Note: The total time Available (Tr) indicates the total usage time for the FlashCopy repository logical drive.

Using the result of these calculations, you should now be able to make an informed decision about whether you should increase the capacity of the FlashCopy repository logical drive. If the repository capacity becomes 100% full during the expected lifetime of the FlashCopy, then you should increase capacity of the FlashCopy repository logical drive. For more information, see “Resizing a FlashCopy repository drive” on page 109.

Using FlashCopy

This section presents an easy-to-follow, step-by-step procedural guide to common administration tasks for a FASiT FlashCopy solution. The tasks that are covered include:

- Checking the status of the Flash Copy premium feature
- Creating a FlashCopy drive
- Mapping a FlashCopy drive to a host
- Viewing FlashCopy drive status
- Disabling a FlashCopy a drive
- Recreating a FlashCopy drive
- Resizing a FlashCopy repository drive
- Deleting a FlashCopy drive

Checking the status of the FlashCopy premium feature


Use this procedure to view a list of premium features on the storage subsystem and to verify that the FlashCopy feature has been enabled. It requires you to:

1. Check the premium options list.
2. View the FlashCopy icon to verify that the feature has been enabled.

Checking the premium option

Perform the following steps to check the premium option:

1. From the Subsystem Management window, click **Storage Subsystem** → **Premium Features** → **List**. Alternately, right-click and click **Premium Features** → **List** (see Figure 22).




SJ000827

Figure 22. Listing premium features

The List Premium Features window opens. It lists the following items:

- The Premium features that are enabled on the storage subsystem
 - The Feature enable identifier
2. Verify that **FlashCopy Logical Drives**: indicates *Enabled* as shown in Figure 23.



SJ000828

Figure 23. Features list showing FlashCopy enabled

Figure 24 shows an example where the FlashCopy copy feature is not enabled.




Figure 24. Feature list showing FlashCopy not enabled

See “Enabling FlashCopy” on page 78 for the procedure to enable FlashCopy.

3. Click **Close** to close the window.

Note: If you receive a *Premium Features - Out of Compliance* error message during a management session, use the Recovery Guru to resolve the problem.

Viewing the FlashCopy icon

Check the status of the FlashCopy feature icon at the bottom left of the device management GUI, as shown in Figure 25.




Figure 25. FlashCopy feature icon

The example in Figure 26 shows a disabled FlashCopy feature.




Figure 26. FlashCopy feature icon disabled

Creating a FlashCopy logical drive

Use this procedure to create a FlashCopy logical drive from a base volume.

Note:

- See Create a FlashCopy Logical Drive Wizard online help for additional instructions for the operating-system-specific procedures. Failure to complete the additional steps that are required for your host operating system can result in an inaccurate point-in-time image of the base logical drive.
- You cannot create a FlashCopy logical drive of a base logical drive that is the secondary logical drive in a Remote Volume Mirror.

Perform the following steps to create a FlashCopy logical drive:

1. Place the application into backup mode:
 - a. Stop the host application that is accessing the base logical drive, sync file system. Unmount the base logical drive if possible. This enables a valid consistent copy to be taken.

It might not be convenient to stop database applications. However, in this case, it is required to place the application into a backup mode or place it in an acquiesced state while you creation the FlashCopy.
 - b. It is important to back up application recovery files, such as role back and redo logs, because these files might be located in different physical disk storage or on different logical drives.
2. Launch the Create a FlashCopy Logical Drive Wizard:
 - a. Select a base logical drive from the Logical view.
 - b. Click **Logical Drive** → **FlashCopy** → **Create**. Alternately, right-click and select **Create FlashCopy Logical Drive**. See Figure 27.




Figure 27. Create FlashCopy menu options

The Create a FlashCopy Logical Drive Wizard opens as shown in Figure 28 on page 89.




Figure 28. Create a FlashCopy Logical Drive Wizard startup

Note: If the FlashCopy logical drive is to be based on the root disk of the host operating system, the final point-in-time image might not be completely consistent with the base logical drive.

3. Create the FlashCopy logical drive:
 - a. Review the information in the initial window as shown in Figure 28. Click **OK** to proceed to the wizard introduction window.

Follow the instructions on each wizard panel, and click the **Next** button when you are ready to continue to the next panel.

Note: Each wizard panel has context-sensitive help. Click the **Help** button on a particular panel to receive help for that panel.

- b. The Introduction window (Figure 29 on page 90) defines what a FlashCopy logical drive is and the physical components that are associated with a FlashCopy logical drive. It enables you to select either the Simple or Advanced path through the Create FlashCopy Logical Drive Wizard:
 - **Simple Path:** Proceeds to the **Specify Name** panel, which provides a preview of the FlashCopy and repository default names. You can also change the defaults on this panel.
 - **Advanced Path:** Proceeds to the **Allocate Capacity** panel. There you select the Free Capacity or Unconfigured Capacity node on which to place the FlashCopy repository logical drive.

If no free capacity exists or the available free capacity is unusable, a warning message displays.





Figure 29. Wizard Introduction window

- c. If you want to use the Advanced path, select **Advanced** and click **Next**. The Allocate Capacity window (Figure 30 on page 91) opens. From this window you can choose free capacity or unconfigured capacity. A description of each option follows.




SJ000835

Figure 30. Allocate Capacity window

- **Free Capacity:**


If you select **Free Capacity** and click **Next**, the Specify Logical Drive Parameters window (Figure 31 on page 92) opens. Define the FlashCopy logical drive-to-LUN mapping parameter and the FlashCopy repository logical drive full conditions. Click **Next**.



SJ000836

Figure 31. Specify Logical Drive Parameters

The Preview window (Figure 32 on page 93) opens and displays the FlashCopy and repository parameters. You can click **Back** to return to the previous windows and edit the parameters or click **Finish** and continue.




SJ000837

Figure 32. Preview window

- **Unconfigured Capacity:**

In the Allocate Capacity window (Figure 30 on page 91), if you select **Unconfigured Capacity** and click **Next**, the Specify Array Parameters window (Figure 33 on page 94) opens. Here you can specify a new array on which the repository resides. You can also specify the RAID level of the array that meets the FlashCopy repository logical drive data storage and protection requirements.



SJ000838

Figure 33. Specify Array Parameters window

After you select where you are placing the volume, same array, and free capacity on another array, or you create a new array, the steps for defining the FlashCopy logical drive parameters are the same as for the Free Capacity option.

On the Specify Logical Drive Parameters window (Figure 31 on page 92), define the FlashCopy logical drive-to-LUN mapping parameter, the threshold percentage full parameter, and the FlashCopy repository logical drive full conditions.

The Preview window (Figure 32 on page 93) provides a preview of the FlashCopy and repository parameters. You can click **Back** and return to the previous windows to edit the repository parameters, or click **Finish** and continue.


- d. On the Specify Names window (Figure 34 on page 95), define the FlashCopy logical drive name and the name of its associated FlashCopy repository logical drive.

The default naming convention for the first FlashCopy uses the base volume name and adds a suffix of "-1" for the logical drive and "-R1" for the repository drive. The second FlashCopy increments this number to 2, and so on, up to the four volumes.

For example, if you are creating the first FlashCopy logical drive for a base logical drive called *DataVol*, then the default FlashCopy logical drive name is *DataVol-1*, and the associated FlashCopy repository logical drive default name is *DataVol-R1*. The default name of the next FlashCopy logical drive you create based on *DataVol* is *DataVol-2*, with the corresponding FlashCopy repository logical drive named *DataVol-R2*, by default.

Change the default names if required.

Click **Next** to continue.



SJ000839

Figure 34. Specify Names window

Tips about names:

- Regardless of whether you use the software-supplied sequence number that (by default) populates the **FlashCopy logical drive name** or **FlashCopy repository logical drive name** field, the next default name for a FlashCopy or FlashCopy repository logical drive still receives a sequence number determined by the software. For example, you might name the first FlashCopy of base logical drive *DataVol* *DataVolMay28*, without using the software-supplied sequence number of 1. Regardless, the software assigns a default name for the next FlashCopy as *DataVol-2*.
- The next available sequence number is based on the number of existing FlashCopies of a base logical drive. If you delete a FlashCopy logical drive, its sequence number becomes available again.
- You must choose a unique name for the FlashCopy and FlashCopy repository logical drives. Otherwise, the system displays an error message.
- Names have a 30-character limit. When you reach this limit in either the **FlashCopy logical drive name** or **FlashCopy repository logical drive name** field, it no longer accepts input. If the base logical drive name contains 30 characters, the default names for the FlashCopy and its associated FlashCopy repository logical drive use the base logical drive name truncated just enough to add the sequence string. For example, for "Host Software Engineering Group GR-1", the default FlashCopy name would be "Host Software Engineering GR-1". The default repository name would be "Host Software Engineering G-R1."

- e. On the Specify FlashCopy Repository Logical Drive's Capacity window (Figure 35), set the repository drive capacity as a percentage of capacity of the base logical drive.



SJ000840

Figure 35. Specify Repository Drive Capacity


Use the **percent (%) of base logical drive** field to set the desired capacity. Click **Next** to continue.

The capacity requirement varies, depending on the frequency and size of I/O writes to the base logical drive and the length of time you need to retain the FlashCopy logical drive. In general, you should choose a larger capacity for the repository if you intend to keep the FlashCopy logical drive for a long period of time. You should choose a larger capacity if a large percentage of data blocks will change on the base logical drive during the life of the FlashCopy logical drive due to a high level of I/O activity. Use historical performance monitor data or other operating system utilities to help you determine typical I/O activity levels on the base logical drive.

Note: In most situations, the 20% default value should be sufficient capacity for your FlashCopy repository logical drive. For information about determining the optimal size, see “Estimating FlashCopy repository logical drive capacity” on page 83, and “Estimating FlashCopy repository life” on page 84.

- f. The Create FlashCopy Logical Drive Wizard - Preview window opens. It displays components that are associated with the FlashCopy. Review the information and click **Next**.


- g. The Create FlashCopy Logical Drive Wizard - Completed window (Figure 36) opens. It displays the associated logical drives and components that make up the FlashCopy relationship. Click **Finish** to continue.



SJ000841

Figure 36. Preview window

The FlashCopy drives are now displayed in the device management GUI, as shown in Figure 37 on page 98. See "Viewing the FlashCopy icon" on page 87 for information about interpreting the display icons.



SJ000842

Figure 37. FlashCopy volumes


- Restart the host application. After you create one or more FlashCopy logical drives, mount the base logical drive and restart the host application using that base logical drive.

Mapping a FlashCopy drive to a host

Assign logical drive-to-LUN mappings between the FlashCopy logical drive and the host that will access the FlashCopy logical drive using the Mappings view of the Subsystem Management window. In some cases, depending on the host operating system and if any logical drive manager software is in use, mapping the same host to both a base logical drive and its associated FlashCopy logical drive can result in conflicts. For operating-system-specific instructions, see “Instructions for Solaris” on page 113 or “Instructions for HP-UX: Logical Drive Manager” on page 116.

Perform the following steps to map the FlashCopy logical drive to a host:


- Open the Mappings view of the Management window. See Figure 38 on page 99. The newly-created FlashCopy logical drive displays in the undefined mapping section.



SJ000843

Figure 38. Undefined FlashCopy disk

2. Select the host or host group that you want to map the drive.
3. Right-click and select **Define Additional Mapping**, as shown in Figure 39.




SJ000844

Figure 39. Define Additional Mapping

4. The Define Additional Mapping window (Figure 40 on page 100) opens.
 - a. Select the FlashCopy drive.
 - b. Select the host or host group.
 - c. Set the LUN number.

d. Click **Add**.




SJ000845

Figure 40. Define Additional Mapping window

In the example shown in Figure 39 on page 99 and Figure 40, the FlashCopy logical disk DataVol-1 is made available to the backup server host group. This enables the backup server to access and mount the disk as part of its own file system. It also allows a local backup to be performed by the backup application software.

It is possible to map the FlashCopy logical disk to the same server that owns the base logical disk. However, immediately after creating the FlashCopy, the two logical disks appear exactly the same (a block-by-block copy). Many operating systems do not tolerate seeing an exact duplicate volume. You might need to complete other steps before you can access it. The mapping is shown in Figure 41 on page 101.

Note: If you use this FlashCopy on a regular basis (for example, for backup purposes), use the Disable FlashCopy and Re-create FlashCopy options to reuse the FlashCopy. Using these options preserves the existing mappings to the FlashCopy logical drive. For more information, see “Disabling a FlashCopy logical drive” on page 106 and “Recreating a FlashCopy logical drive” on page 108.



SJ000846

Figure 41. Mapped FlashCopy logical disk

Finally, use specific operating-system and host utilities to mount and use the mapped FlashCopy drive.

1. Run the host-based hot_add utility to register the FlashCopy logical drive with the host operating system.
2. Run the host-based SMdevices utility to associate the mapping between the physical device name and the logical drive name.
3. Mount the logical drive to the host.

For information about specific host operating-system procedures, see “Instructions for Solaris” on page 113 or “Instructions for HP-UX: Logical Drive Manager” on page 116.

Viewing the FlashCopy logical drive status

You can determine the status of the FlashCopy logical drive by viewing the icons that change depending of the state of the drive. The logical drive component property display is also useful in determining the state of the logical drives.

Use the FlashCopy Repository Logical Drive - Properties window to view the FlashCopy repository logical drive base and capacity properties. You can also use this window to specify the capacity percentage full and the action to be taken if the FlashCopy repository logical drive becomes full.

The progress of modification operations displays at the bottom of the window.

FlashCopy icon states

To view the FlashCopy icon, open the Storage Management Device Manager GUI Physical/Logical view. The icon states are shown in Figure 42 on page 102.

FlashCopy status representation in logical view		
<u>Logical drive status icon</u>	<u>Mirror status</u>	<u>Icon</u>
FlashCopy	Optimal	
	Disabled	
	Failed	
	Offline	
Repository	Optimal	
	Degraded	
	Failed	
	Full	
	Offline	
	Warning	
Free capacity	Spare capacity	


SJ000847

Figure 42. FlashCopy icon states

FlashCopy repository properties

Perform the following steps to view the FlashCopy repository properties:


1. Select a FlashCopy repository logical drive in the Logical view of the Subsystem Management window.
2. Click **Logical Drive** → **Properties**, or right-click and select **Properties**, as shown in Figure 43 on page 103.



SJ000848

Figure 43. Repository Logical Drive Properties

The FlashCopy Repository Logical Drive - Properties window opens, as shown in Figure 44 on page 104.




SJ000849

Figure 44. Base Repository Logical Drive Properties

The **Base** tab of the FlashCopy Repository Logical Drive - Properties window displays the following information for the selected FlashCopy repository logical drive:

- Logical drive name
 - Worldwide name
 - Status
 - Controller ownership
 - Capacity
 - RAID level
 - Modification priority
 - Associated base logical drive
 - Associated FlashCopy logical drive
3. Click the **Capacity** tab (Figure 45 on page 105) to view or set the following FlashCopy repository logical drive capacity properties:
- **FlashCopy repository logical drive percentage full**
Allows a threshold level to be set for the FlashCopy repository logical drive capacity. When the defined percentage is reached, a warning is issued indicating that the repository is nearing its capacity. The default percentage setting is 50% of the maximum capacity of the FlashCopy repository logical drives.
Use the **percent (%) full** field to define the percentage at which a warning is issued.



SJ000850

Figure 45. Repository drive warning threshold

- **Repository full policy**

When a FlashCopy repository logical drive reaches its capacity and becomes full, one of the following actions occurs:

- Fail FlashCopy logical drive
- Fail writes to base logical drive

If the FlashCopy repository logical drive is set to *Fail FlashCopy logical drive* (as shown in Figure 45) when it becomes full, its data is not recoverable and the FlashCopy cannot be accessed. In this case, the only available option is to delete the FlashCopy logical drive or recreate the FlashCopy logical drive to create a new point-in-time image.

If the FlashCopy repository logical drive is set to *Fail writes to the base logical drive*, the data is recoverable. However, the FlashCopy repository logical drive capacity must be increased so writes to the base logical drive are not rejected. For more information, see “Resizing a FlashCopy repository drive” on page 109.

Viewing the progress of a modification operation

The progress bar at the bottom of the FlashCopy Repository Logical Drive Properties window displays the progress of an operation. You can view the progress of the following operations:

- Copyback
- Reconstruction
- Initialization
- RAID level change

- Dynamic logical drive expansion
- Capacity increase
- Defragmentation
- Segment size change

Note: The storage management software cannot obtain progress information from the storage subsystem controllers if the network management connection to the controllers is down or if the storage subsystem is partially managed. For more information about a partially managed storage subsystem or an unresponsive controller or storage subsystem condition, see the Enterprise Management window online help.

Disabling a FlashCopy logical drive

If you no longer need a FlashCopy logical drive, you might want to disable it. As long as a FlashCopy logical drive is enabled, your storage subsystem performance can be impacted by copy-on-write activity that is directed to the associated FlashCopy repository logical drive. The copy-on-write activity concludes when you disable a FlashCopy logical drive.

If you disable a FlashCopy logical drive instead of deleting it, you can retain it and its associated repository for future use. When you need to create a different FlashCopy of the same base logical drive, you can use the Recreate option to reenable the previously disabled FlashCopy. This takes less time than creating a new FlashCopy.

When you disable a FlashCopy logical drive, note the following points:

- You cannot use that FlashCopy logical drive again until you use the Recreate option on that logical drive.
- Only that FlashCopy logical drive is disabled. All other FlashCopy logical drives remain functional.

If you do not intend to recreate a FlashCopy, you can delete that FlashCopy logical drive instead of disabling it.

Perform the following steps to disable a FlashCopy logical drive:

1. Select the FlashCopy logical drive. Right-click and select **Disable**, as shown in Figure 46 on page 107.




Figure 46. Choosing to disable the FlashCopy drive

2. The Disable FlashCopy Logical Drive confirmation window opens. See Figure 47. On this window, type Yes and click **OK** to begin the disable operation.




Figure 47. Disable FlashCopy confirmation window

The FlashCopy icon in the Physical/Logical View now displays as disabled, as shown in Figure 48 on page 108.




Figure 48. Icon showing the disabled FlashCopy logical drive

Recreating a FlashCopy logical drive

Recreating a FlashCopy logical drive takes less time than creating a new one. If you have a FlashCopy logical drive that you no longer need, instead of deleting it, you can reuse it (and its associated FlashCopy repository logical drive) to create a different FlashCopy logical drive of the same base logical drive.


When you recreate a FlashCopy logical drive, note the following points:

- The FlashCopy logical drive must be in either an optimal or a disabled state.
- All copy-on-write data on the FlashCopy repository logical drive is deleted.
- FlashCopy and FlashCopy repository logical drive parameters remain the same as the previously disabled FlashCopy logical drive and its associated FlashCopy repository logical drive. After you recreate a FlashCopy logical drive, you can change parameters on its associated FlashCopy repository logical drive through the appropriate menu options.
- The system retains the original names for the FlashCopy and FlashCopy repository logical drives. You can change these names, however, after the Recreate option completes.
- When using this option, the previously-configured FlashCopy name, parameters, and FlashCopy repository logical drive are used.

Perform the following steps to recreate a FlashCopy drive:

1. Select the FlashCopy logical drive. Right-click and select **Re-Create**, as shown in Figure 49 on page 109.
2. The Re-create FlashCopy Logical Drive window opens. Type Yes and click **OK**.

This command disables and recreates the FlashCopy logical drive (if it had not been previously disabled) and displays in the Logical view in an Optimal state. The creation timestamp shown on the FlashCopy Logical Drive Properties window is updated to reflect the new point-in-time image. Copy-on-write activity resumes to the associated FlashCopy repository logical drive.



SJ000854

Figure 49. Recreating a FlashCopy

Note:

- To use the Recreate option, the FlashCopy logical drive must be in either an optimal state or a disabled state.
- If the FlashCopy logical drive is in an optimal state, the process first disables and then recreates the FlashCopy logical drive. This process invalidates the current FlashCopy.

Resizing a FlashCopy repository drive

Use this option to increase the storage capacity of an existing FlashCopy repository logical drive. Typically, this option is used when a warning is received that the FlashCopy repository logical drive is in danger of becoming full.

You can achieve an increase in storage capacity by:

- Using free capacity available on the array of the FlashCopy repository logical drive.
- Adding unconfigured capacity (in the form of unused drives) to the array of the FlashCopy repository logical drive. Use this option when no free capacity exists on the array.

Note: A maximum of two drives can be added at one time to increase FlashCopy repository logical drive capacity.

The storage capacity of a FlashCopy repository logical drive cannot be increased if:

- One or more hot spare drives are in use in the logical drive.
- The logical drive has a non-optimal status.
- Any logical drive in the array is in any state of modification.
- The controller that owns this logical drive is in the process of adding capacity to another logical drive. Each controller can add capacity to only one logical drive at a time.
- No free capacity exists in the array.

- No unconfigured capacity (in the form of drives) is available to add to the array.

Perform the following steps to resize a FlashCopy repository drive:

1. Select a FlashCopy repository logical drive in the Logical view of the Subsystem Management window.
2. Click **Logical Drive** → **Increase Capacity**, or right-click and select **Increase Capacity**.

Note: If no free capacity or unconfigured capacity is available, the Increase Capacity option is not available.

The Increase Repository Capacity window opens. You can see the FlashCopy repository logical drive name, the associated FlashCopy logical drive name, the associated base logical drive name, the current capacity, and the amount of free capacity that is available for the selected repository.

If free capacity is available, the maximum free space is shown in the **Increase Capacity by** field. If there is no free capacity on the array, the free space that is shown in the **Increase Capacity by** field is 0. Add drives to create free capacity on the array of the standard logical drive.

3. Use one of the following two methods to increase capacity:
 - Increase FlashCopy repository logical drive capacity using free capacity on the array of the FlashCopy repository logical drive:
 - a. Accept the final capacity increase or use the **Increase Capacity by** field to adjust the capacity. Click **OK**.
 - b. A confirmation window displays. Type Yes and click **OK** to continue.
The Logical view is updated. The FlashCopy repository logical drive with its capacity increased shows a status of Operation in Progress, together with its original capacity and the total capacity being added.
 - c. In addition, the Free Capacity node involved shows a reduction in capacity. If all of the free capacity is used to increase the logical drive size, then the Free Capacity node involved is removed from the Logical view.
 - Increase FlashCopy repository logical drive capacity by adding unconfigured capacity (drives) to the array of the FlashCopy repository logical drive:
 - a. If no unassigned drives are available and empty slots in the drive enclosures are available, insert new drives.
If no unassigned drives are available, and there are no empty slots available in the drive enclosures, install another drive enclosure and additional drives.
 - b. Select **Add Drives**.
The Increase Repository Capacity - Add Free Capacity window opens. Enclosure, slot, and usable capacity details for the available free drives are displayed.

Note: The drives that are displayed have a capacity that is either equivalent to or larger than those that are already employed by the array.
 - c. Select a single drive, or two drives, to be added:
 - Press **Ctrl** and click **Enter** to select the nonadjacent drives.
 - Press **Shift** and click **Enter** to select the adjacent drives.
 - d. Click **Add**.

The Add Free Capacity window closes. Check the **Drives to add [enclosure, slot]** field to ensure that the correct drives are added.

- e. Accept the final capacity or use the **Increase Capacity by** field to adjust the capacity.
- f. Click **OK**.
- g. A confirmation window opens. Type Yes to confirm the operation, and click **OK** to continue.

The Logical view is updated. The FlashCopy repository logical drive having its capacity increased shows a status of *Operation in Progress*, together with its original capacity and the total capacity being added. In addition, the Free Capacity node involved in the increase shows a reduction in capacity.

If all of the free capacity is used to increase the logical drives size, then the Free Capacity node involved is removed from the Logical View.

If a Free Capacity node did not exist before you added capacity and not all of the capacity that is added will be used to increase the FlashCopy repository logical drives capacity, a new Free Capacity node is created and displayed in the Logical view.

Unassigned drives (unconfigured capacity) that are added to increase the FlashCopy repository logical drives capacity change in the Physical view to assigned drives, and become associated to the array of the FlashCopy repository logical drive.

4. View the progress of the capacity increase process. Select the FlashCopy repository logical drive. Click **Logical Drive** → **Properties**, or right-click and select **Properties**.

The FlashCopy Repository Logical Drive - Properties window opens. A progress bar at the bottom of the window indicates the status of the capacity increase.

Deleting a FlashCopy drive

Use this option to delete a FlashCopy logical drive that is no longer needed for backup or application testing purposes. This option results in an increase of free capacity in the array or additional unconfigured capacity.

Note:

- Deleting a logical drive causes loss of all data on the logical drive. Back up the data and stop all I/O before performing this operation, if necessary.
- If a file system is mounted on the logical drive, unmount it before you performing this operation.
- Deleting a base logical drive automatically deletes the associated FlashCopy logical drive and FlashCopy repository logical drive.
- Deleting a FlashCopy logical drive automatically deletes the associated FlashCopy repository logical drive.
- You cannot delete a FlashCopy repository logical drive using the Delete Logical Drive option, but you can by using the **delete** command in the script editor or CLI. See the Enterprise Management window online help system for more information about using the **delete** command.

Perform the following steps to delete a FlashCopy drive:

1. Select the FlashCopy logical drive in the Logical view.

2. Click **Logical Drive** → **Delete**, or right-click and select **Delete**, as shown in Figure 50.




Figure 50. Deleting the FlashCopy logical drive

3. The Delete FlashCopy Logical Drive window opens, as shown in Figure 51. Type **Yes** and click **OK**.




Figure 51. Delete FlashCopy Logical Drive Dialog window

The FlashCopy logical drive and FlashCopy repository logical drive are deleted, destroying any data that they might have contained.

Using FlashCopy with UNIX regular disks

This section explains the procedures for using FlashCopy with UNIX regular disks. It outlines the key steps required to create and reuse a FlashCopy logical drive on the following UNIX-based operating systems, using regular disks:

- Solaris 2.6, 2.7 (Solaris 7), and 2.8 (Solaris 8)
- HP-UX 11.0 (or higher)

Instructions for Solaris

Use the following procedures to create and reuse a FlashCopy logical drive using regular disks on Solaris 2.6, 2.7 (Solaris 7), and 2.8 (Solaris 8). Failure to complete the steps listed may result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives can either be created for reuse (frequent or nightly backups) or for one-time use (speculative change or upgrade testing).

Creating a FlashCopy logical drive

Before you create a FlashCopy logical drive, consider the following questions:

- Are you creating the FlashCopy logical drive for immediate use?
If you plan to use the FlashCopy logical drive immediately,, stop all I/O activity to the base logical drive. Also, unmount the base logical drives from the host to which they are currently mounted. This ensures that it will capture an accurate point-in-time image of the data on the base logical drive.
- Are you creating the FlashCopy logical drive for use at a later date?
If you plan to use the FlashCopy logical drive at a later date, do not unmount or stop I/O activity to the base logical drive at this time. You must perform these activities immediately before you use the FlashCopy logical drive for the first time.

Perform the following steps to create a FlashCopy logical drive:

1. Start the Storage Management software. The Enterprise Management window opens.
2. Within the Logical view of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive using one of the following methods:
 - Use the Create FlashCopy Logical Drive Wizard. This is accessed using a Subsystem Management window.
 - Create a FlashCopy logical drive using either the script editor or a command-line shell on the host operating system. Type the following command and press Enter.

```
create FlashCopylogicaldrive
```

3. Within the Logical view of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive using one of the following methods:
 - Use the Create FlashCopy Logical Drive Wizard. This is accessed using a Subsystem Management window.

- Create a FlashCopy logical drive using either the script editor or a command-line shell on the host operating system. Type the following command and press Enter.

```
create FlashCopylogicaldrive
```

4. If supported by the operating system, run the hot_add utility (or operating system-specific utility), or restart the host on which the FlashCopy will be used. This ensures that the host operating system recognizes the FlashCopy logical drive.

After you have created logical drives and defined logical drive-to-LUN mappings, run the hot_add utility to ensure that the operating system is aware of the newly-created logical drives, without having to restart the host. For information about which operating systems support the hot_add utility, see the Storage Manager Software Installation Guide for your specific operating system, which is shipped on a CD with the product. You can also download it from the Web site at:

<http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700>

For the Solaris operating system, at the host prompt, type the following command and press Enter.

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

After logical drives are created and logical drive-to-LUN mappings are defined, this step ensures that the operating system is aware of the newly-created logical drives, without having to restart the host.

5. Run the SMdevices utility to associate the LUN with a host operating system device and to ensure that the FlashCopy logical drive is recognized by the host.

After you have created logical drives and defined logical drive-to-LUN mappings, run the SMdevices utility to ensure that the logical drive name and the operating system device name (assigned by the operating system) correlate.

6. If you plan to use the FlashCopy logical drive immediately, go to step 7. If you plan to use the FlashCopy logical drive at a later date, disable the FlashCopy logical drive now using one of the following methods:

- Click **Logical Drive**—> **FlashCopy** —> **Disable** in the Subsystem Management window.
- Disable a FlashCopy logical drive using the script editor or a command-line shell on the host operating system. Type the following command and press Enter:

```
disable FlashCopy logical drive
```

Note: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive or re-enable data transfer at this time.

7. Mount the FlashCopy logical drive to its intended host.
8. Use the FlashCopy logical drive with your backup application, for speculative testing, or with another application.
9. Unmount the FlashCopy logical drive.

10. After the FlashCopy logical drive is no longer needed, disable or delete the FlashCopy logical drive.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can recreate the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and avoids any reduction in performance that might occur if the FlashCopy logical drive remains available.

Reusing FlashCopy logical drives

Typically, After a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is needed.

Perform the following steps to create a new point-in-time image of the same base logical drive, complete the following steps:

Note: Stop all I/O activity to the base logical drive at this point (or suspend data transfer). This ensures that an accurate point-in-time image of the base logical drive is captured.

1. Unmount the base logical drive.
2. In the Storage Management software, recreate the FlashCopy logical drive using one of the following methods:
 - Click **Logical Drive** → **FlashCopy** → **Re-create** from the menus in the Subsystem Management window.
 - Recreate a FlashCopy logical drive using the script editor or a command-line shell on the host operating system. Type the following command and press Enter.

```
create FlashCopylogicaldrive
```

3. Remount the base logical drive to its original host.
4. Mount the FlashCopy logical drive to its intended host.

Note: If you stopped I/O activity to the base logical drive or suspended data transfer, resume I/O activity to the base logical drive or re-enable data transfer at this time.

5. Use the FlashCopy logical drive with your backup application or with another application.
6. Unmount the FlashCopy logical drive.
7. When the FlashCopy logical drive is no longer needed, disable the FlashCopy logical drive.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can recreate the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that might occur if the FlashCopy logical drive remains available.

Instructions for HP-UX: Logical Drive Manager

Use the following procedures to create FlashCopy logical drives on a host running HP-UX 11.0 (or higher) using Logical Volume Manager (LVM). Failure to complete the steps listed can result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives can either be created for repeated reuse (frequent or nightly backups) or for one-time use (speculative change or upgrade testing).

Creating a FlashCopy logical drive

Perform the following steps to create a FlashCopy logical drive:

1. Start the Storage Management software. The Enterprise Management window opens.
2. Launch a Subsystem Management window, using one of the following methods:
 - Select the storage subsystem in the Device Tree view or from the Device Table. Then, click the **Manage Device** toolbar button, or click **Tools** → **Manage Device**.
 - Right-click the storage subsystem in the Device Tree View or Device Table and select **Manage Device**.
 - Double-click a storage subsystem node in the Device Table.
 - Select the storage subsystem in either the Device Tree view or from the Device Table and press Enter.

The Subsystem Management window opens in a separate window.

Note: Stop the host application that is accessing the base logical drive and unmount the base logical drive. Unmounting the base logical drive does not apply when the base logical drive is the root disk of the host operating system.

3. When creating a FlashCopy logical drive based on a mounted file system, always perform a sync to flush the file system cache immediately prior to creating a FlashCopy logical drive.

At the host prompt, type the following command and press Enter.

```
sync
```

All unwritten file system buffers are flushed.

4. Within the Logical view of the Subsystem Management window, select a standard logical drive, and create a FlashCopy logical drive using one of the following methods:
 - Use the Create a FlashCopy Logical Drive Wizard. This is accessed using an Subsystem Management window.
 - Create a FlashCopy logical drive using either the script editor or a command-line shell on the host operating system. Type the following command and press Enter.

```
create FlashCopylogicaldrive
```

5. Assign a logical drive-to-LUN mapping between the FlashCopy logical drive and the host that will access the FlashCopy logical drive. Logical drive-to-LUN mappings can be defined using one of the following methods:

- Use the Storage Partitioning Wizard to quickly define a single storage partition. It guides you through the major steps to specify which host will access a logical drive, and the associated LUNs.
- Create a logical drive-to-LUN mapping using the script editor or a command-line shell on the host operating system. Type the following command and press Enter.

```
create mapping logical drive
```

6. At the host prompt, type the following command and press Enter.

```
ioscan -fn
```

A list of the mapped devices recognized by the host is displayed.

Note: If the required device names are not displayed using this command, at the host prompt, type the following command and then press Enter.

```
insf
```

7. Remount the base logical drive to its original host.

Note: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive or reenable data transfer at this time.

8. Perform the following steps to import the FlashCopy logical drives into the Logical Drive Manager:
 - a. Create a new directory for the new array, as shown in the following example.

```
mkdir /dev/vg02
```

- b. Create a group node for the new array, as shown in the following example.

```
mknod /dev/vg02/group c 64 -0x020000
```

- c. Import the FlashCopy logical drive LUNs. At the host prompt, type the following command and press Enter.

```
vgimport /dev/vg02 FlashCopy-block-node-1 FlashCopy-block-node-2
```

The following example addresses how to handle redundancy with multiple paths. The two paths or nodes represent the primary and alternate paths of the FlashCopy volume. Please consider the following example:

```
vgimport /dev/vg02 /dev/dsk/c66t0d1 /dev/dsk/c69t0d1
```

Note: You must verify that the `/dev/dsk` device files exist on the FlashCopy logical drive. Use the `SMdevices` utility or the HP-UX `ioscan` utility to perform the verification.

The system displays a warning indicating that a backup of the array being imported might not exist on the host. This message is only a warning and requires no response. The import continues and completes successfully.

The backup for this array is created later, when it is exported.

d. Activate the new array as shown in the following example:

```
vgchange -a y /dev/vg02
```

9. If a file system existed on the base logical drive, then it also exists on the FlashCopy logical drive. Before you mount the FlashCopy logical drive, however, perform a file system check to ensure that the file system is consistent, for example:

```
fscck /dev/vg02/1vol01
```

10. Mount the FlashCopy logical drive to its intended host.
11. Use the FlashCopy logical drive with your backup application, for speculative testing, or with another application.
12. Unmount the FlashCopy logical drive.
13. Disable the FlashCopy logical drive using one of the following methods:
 - Click **Logical Drive** → **FlashCopy** → **Disable** in the Subsystem Management window.
 - Disable a FlashCopy logical drive using the Script Editor or a command line shell on the host operating system. Type the following command and press Enter.

```
disable FlashCopylogicaldrive
```

Disable the FlashCopy logical drive after it is no longer needed.

If you disable a FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can recreate the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and avoids any reduction in performance that might result by retaining the availability of an unused FlashCopy logical drive.

Reusing FlashCopy logical drives

Typically, after a FlashCopy logical drive is created it is disabled until a new point-in-time image of the same base logical drive is needed. Perform the following steps to create a new point-in-time image of the same base logical drive:

Note: Stop all I/O activity to the base logical drive (or suspend data transfer). This ensures that an accurate point-in-time image of the base logical drive is captured.

To create a FlashCopy logical drive, complete the following steps:

1. Start the Storage Management software. The Enterprise Management window opens.
2. Launch a Subsystem Management window using one of the following methods:

- Select the storage subsystem in either the Device Tree View or Device Table. Then, click the Manage Device toolbar button or click **Tools—> Manage Device**.
- Right-click the storage subsystem in the Device Tree View or Device Table, and select **Manage Device**.
- Double-click a storage subsystem node in the Device Table.
- Select the storage subsystem in the Device Tree View or the Device Table, and then press **Enter**.

The Subsystem Management window opens in a separate window.

3. Unmount the base logical drive.
4. In the Storage Management software, recreate the FlashCopy logical drive using one of the following methods:
 - Click **Logical Drive —> FlashCopy —> Re-create** in the Subsystem Management window.
 - Recreate a FlashCopy logical drive using the script editor or a command-line shell on the host operating system. Type the following command and press Enter.

```
recreateFlashCopy logical drive
```

5. Remount the base logical drive to its original host.

Note: If you stopped I/O activity to the base logical drive or suspended a data transfer, resume I/O activity to the base logical drive or re-enable data transfer at this time.

6. Perform the following steps to import the FlashCopy logical drives into the Logical Drive Manager:
 - a. Create a new directory for the new array as shown in the following example.

```
mkdir /dev/vg02
```

- b. Create a group node for the new array as shown in the following example.

```
mkknod /dev/vg02/group c 64 -0x020000
```

- c. Import the FlashCopy logical drive LUNs. At the host prompt, type the following command and press Enter.

```
vgimport /dev/vg02 FlashCopy-block-node-1 FlashCopy-block-node-2 ...
```

An alternate command can be used as shown in the following example.

```
vgimport /dev/vg02 /dev/dsk/c66t0d1 /dev/dsk/c69t0d1
```

Note: You must verify that the `/dev/dsk` device files must exist on the FlashCopy logical drive. Use the `SMdevices` utility or the HP-UX `ioscan` utility to perform this verification.

A warning is displayed indicating that a backup of the array being imported might not exist on the host. This message is only a warning and can be ignored. The import continues and completes successfully.

The backup for this array is created when it is exported later.

- d. Activate the new array, as shown in the following example.

```
vgchange -a y /dev/vg02
```

7. If a file system existed on the base logical drive, then it also exists on the FlashCopy logical drive. However, before the FlashCopy logical drive can be mounted, submit it to a file system check.

Perform a file system check to ensure that the file system is consistent, for example:

```
fsck /dev/vg02/lvo101
```

8. Mount the FlashCopy logical drive to its intended host.
9. Use the FlashCopy logical drive with your backup application or other application.
10. Unmount the FlashCopy logical drive.
11. If the FlashCopy logical drive is no longer needed, disable the FlashCopy logical drive.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can recreate the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that might occur if the FlashCopy logical drive remains available.

Script editor and command-line interface

Many storage management options that are available through the Subsystem Management window can be sent to the storage subsystem using statements in scripts. You can use the script editor to create or edit a script file, save a script file to the local disk of the Storage Management station, or load a script file from disk. You can also use the command-line interface (CLI) to issue individual commands to the scripting engine from the host operating system command shell or to call complete prewritten scripts.

See the *IBM TotalStorage FAST700 and Copy Services Redbook*, SG24–6808 for information about using the script editor, examples of automated scripts, and information about the command-line interface (CLI). You will find this redbook at:

www.ibm.com/redbooks

Missing logical drives

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

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

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

If missing logical drives are detected by the storage subsystem, a Missing Logical Drives group is created in the Logical view of the Subsystem Management window. Each missing logical drive is shown and identified by its worldwide name and logical drive type. Missing logical drives are identified as being either a standard logical drive, base logical drive, FlashCopy repository logical drive, primary logical drive, secondary logical drive, or mirror repository logical drive.

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

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

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

Appendix A. AIX system requirements

This appendix lists the minimum hardware and software requirements that an AIX system must meet to be used in a FASTt200, FASTt500, or FASTt700 storage subsystem that is managed by Storage Manager 8.21. A README Text file containing the latest installation and user information about the storage management software, AIX filesets, and hardware components are located in the AIX subdirectory on the FASTt Installation CD. The most recent copy, which supersedes this document, is maintained on the following Web site:

<http://www.pc.ibm.com/qtechinfo/MIGR-43793.html>

Once there, follow the link to v8.21 information.

Hardware requirements

The following hardware is required at a minimum.

- Any RS/6000® or eServer pSeries™ server that supports FC 6227 or FC 6228 host bus adapters
- Two host bus adapters (HBAs)

Software requirements

The following software must be at the specified versions.

- AIX 4.3.3 or AIX 5.1
- Filesets associated with IBM RDAC driver installations listed in Table 10 or Table 11

RDAC installation requirements

The following filesets must be at the specified versions or later.

- For AIX 4.3.3:

Table 10. 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	4.3.3.83
devices.common.IBM.fc.rte	4.3.3.75
devices.pci.df1000f7.com	4.3.3.83
devices.pci.df1000f7.rte	4.3.3.75
devices.pci.df1000f9.rte	4.3.3.75
devices.scsi.sarray.rte	4.3.3.50

- For AIX 5.1:

Table 11. Filesets required for AIX 5.1 RDAC

PTF filesets	Version
devices.fcp.disk.array.diag	5.1.0.25
devices.fcp.disk.array.rte	5.1.0.37

Table 11. Filesets required for AIX 5.1 RDAC (continued)

PTF filesets	Version
devices.common.IBM.fc.rte	5.1.0.10
devices.pci.df1000f7.com	5.1.0.28
devices.pci.df1000f7.rte	5.1.0.15
devices.pci.df1000f9.rte	5.1.0.15
devices.scsi.sccarray.rte	5.1.0.0

Note: The AIX RDAC driver files are not included on the FAStT installation CD. Either install them from the AIX Operating Systems CD, if the correct version is included, or download them directly from the Web site as follows::

techsupport.services.ibm.com/server/fixes

For downloading instructions, see “Installing host RDAC software on AIX hosts” on page 14.

Appendix B. HP-UX system requirements

This appendix lists the minimum hardware and software requirements that an HP-UX system must meet to be used in a FASTt200, FASTt500, or FASTt700 storage subsystem that is managed by Storage Manager 8.21.

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 a SMruntime environment

Ensure that the HP-UX host is running one of the following operating systems with the appropriate patches. Because patches can be superseded by more recent versions, refer to your operating system documentation or contact your operating system supplier to ensure that you have the correct patches.

- HP-UX 11.0 (32 bit or 64 bit) with the following patches:
 - B.11.00.53.2
 - B.11.00.53.8
 - PHKL_23939, version 1.0
 - PHKL_24004, version 1.0
 - PHKL_24027, version 1.0
 - PHKL_24612, 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 11.11 and HP-UX 11i with the following patches:
 - B.11.11.09
 - PHKL_25304
 - PHKL_25375
 - PHKL_25896
 - PHKL_PHLK_26519
- 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 12.

Table 12. HP-UX kernel parameter configuration requirements

Parameter	Description	Configuration
max_thread_proc 64	Maximum threads per process	1024
maxfiles	Soft-file limit per process	2048
maxuser	Influences other parameters	256 or greater
ncallout	Number of pending timeouts	4144

Appendix C. Solaris system requirements

This appendix lists the minimum hardware and software requirements that a Solaris system must meet to be used in a FAStT200, FAStT500, or FAStT700 storage subsystem that is managed with Storage Manager 8.21.

The Solaris system must be an Sparc S20 processor with:

- 256 MB system memory
- CD-ROM drive
- Mouse or similar pointing device
- Ethernet network interface card
- 1 MB available on /opt and root (or root-equivalent) privileges for installing RDAC

Ensure that the Solaris host is running one of the following operating systems with the appropriate patches. Because patches can be superseded by more recent versions, refer to your operating system documentation or contact your operating system supplier to ensure that you have the correct patches.

- Solaris 2.6 with the following patches:
 - 105181-32
 - 105356-19
 - 106429-02
- Solaris 7 with the following patches:
 - 106541-21
 - 107544-03
 - 108376-38
- Solaris 8 with the following patches:
 - 108528-15
 - 111111-03
 - 111310-01
 - 111293-04
- 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>
```

This precaution applies only when you use the `-s` option to import devices or disks with redundant paths. The condition occurs because the `-s` option causes a search on the system for each disk when used with **vgimport**.

When the new list is prepared, the links that were previously designated as primary and alternate might not remain as they had been configured on node *A*.

Use the following procedure to correct the paths that were changed in this way when using **vgimport -m -s** with LVM commands.

Note: The following instructions are for a two-node cluster. You can easily expand or adapt these instructions to accommodate a cluster which has more than two nodes.

1. Identify and note the primary and alternate paths for each device by typing the following command:

```
# SMdevices
```

2. Compare the primary and alternate paths for each device on node *A* to Node *B* by using the **vgdisplay -v** command. As previously noted, an expected behavior of the **vgimport -m -s** command is that the primary and alternate paths change during export.

On node *A*:

```
#vgdisplay -v <volume_group_name>
```

On node *B*:

```
#vgdisplay -v <volume_group_name>
```

3. If the original primary path of a disk becomes an alternate path after the newly imported volume group entry is created in `/etc/lvmtab`, the order can be easily reverted by using **vgreduce** to remove the wrong primary path and then using **vgextend** to add the path back again:

```
#vgreduce  
vg1 /dev/dsk/<device_name>#vgextend vg1  
/dev/dsk/<device_name
```

Appendix E. JNI Host Bus Adapter settings

The following table detail settings for the various host bus adapter (HBA) cards for Sun Solaris. These cards are not plug-and-play with auto-configuration. Instead, they may require manual alteration of the settings/bindings.

Table 13. Configuration file name: /kernel/drv/fca-pci.conf

Original value	New value	Comments
fca_nport = 0;	fca_nport = 1;	
ip_disable = 0;	ip_disable = 1;	
failover = 0;	failover = 30;	
busy_retry_delay= 5000;	busy_retry_delay = 5000;	
link_recovery_delay = 1000;	link_recovery_delay = 1000;	
scsi_probe_delay = 5000;	scsi_probe_delay = 5000;	
def_hba_binding = "fca-pci*";	def_hba_binding = "nonjni";	
def_wwnn_binding = "\$xxxxxx"	def_wwnn_binding = "xxxxxx"	Uncomment the line.
def_wwpn_binding = "\$xxxxxx"	same	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 14. Configuration file name: /kernel/drv/jnic146x.conf

Original value	New value	Comments
FcLoopEnabled = 1	FcLoopEnabled = 0	Uncomment the line.
FcFabricEnabled = 0	FcFabricEnabled = 1	Uncomment the line.
FailoverDelay = 30;	FailoverDelay = 30;	Uncomment the line.
JniCreationDelay = 5;	JniCreationDelay = 10;	Uncomment the line.
def_wwnn_binding = "\$xxxxxxxxxxx"	def_wwnn_binding = "xxxxxxxxxxx"	Uncomment the line.
def_wwpn_binding = "\$xxxxxxxxxxx"	same as the original entry	Uncomment the line.

Table 14. Configuration file name: /kernel/drv/jnic146x.conf (continued)

Original value	New value	Comments
Add	target0_hba = "jnic146x0";	
Add	target0_wwpn = "<controller wwpn>"	
Add	target1_hba = "jnic146x1";	
Add	target1_wwpn = "<controller wwpn>"	

Note: You might need to run the /etc/raid/bin/genjnicconf reconfigure script from the Solaris shell.

```
# /etc/raid/bin/genjnicconf
```

Table 15. Configuration file name: /kernel/drv/jnic.conf

Original value	New value	Comment
FcLoopEnabled = 1	FcLoopEnabled = 0	Uncomment the line.
FcFabricEnabled = 0	FcFabricEnabled = 1	Uncomment the line.
FailoverDelay = 30;	FailoverDelay = 30;	Uncomment the line.
JniCreationDelay = 5;	JniCreationDelay = 10;	Uncomment the line.
def_wwnn_binding = "\$xxxxxxxxxxx"	def_wwnn_binding = "xxxxxxxxxxx"	Uncomment the line.
def_wwpn_binding = "\$xxxxxxxxxxx"	same as the original entry	Uncomment the line.
Add	target0_hba = "jnic146x0";	
Add	target0_wwpn = "<controller wwpn>"	
Add	target1_hba = "jnic146x1";	
Add	target1_wwpn = "<controller wwpn>"	

Note:

- You might need to run the /etc/raid/bin/genjnicconf reconfigure script from the Solaris shell.

```
# /etc/raid/bin/genjnicconf
```

- Set *portEnabled = 1*; only when you see JNI cards entering non-participating mode in the /var/adm/messages file. Under that condition:

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 16. Configuration file name: /kernel/drv/fcaw.conf

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

Table 16. Configuration file name: /kernel/drv/fcaw.conf (continued)

Original value	New value	Comments
ip_disable = 0;	ip_disable=1;	
failover = 0;	failover =30;	
busy_retry_delay = 5000;	busy_retry_delay = 5000;	
link_recovery_delay = 1000;	link_recovery_delay = 1000;	
scsi_probe_delay = 5000;	scsi_probe_delay = 5000;	
def_hba_binding = "fcaw*";		Uncomment the line.
def_wwnn_binding = "\$xxxxxx"	def_wwnn_bindindef_hba_binding = "nonjni"; g = "xxxxxx"	Uncomment the line.
def_wwnn_binding = "\$xxxxxx"	Same as the original entry.	Uncomment the line.
Will be added by reconfigure script	name="fcaw" parent="<physical path>"unit-address="<#>"	
Will be added by reconfigure script	target0_hba="fcaw0" target0_wwpn="<controller wwpn>";	
Will be added by reconfigure script	name="fcaw" parent="<physical path>"unit-address="<#>"	
Will be added by reconfigure script	target0_hba="fcaw0" target0_wwpn="<controller wwpn>";	

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

```
# /etc/raid/bin/gen SCSIconf
```

Notices

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publications. IBM may make improvements or changes (or both) in the product(s) or program(s) (or both), described in this publication at any time without notice.

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

Some software may differ from its retail version (if available), and may not include user manuals or all program functionality.

Trademarks

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

IBM
AIX
e(logo)

Netfinity
RS/6000
TotalStorage

Microsoft Windows and Windows NT are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

IBM agreement for licensed internal code

You accept the terms of this Agreement¹ by your initial use of a machine that contains IBM Licensed Internal Code (called "Code"). These terms apply to Code used by certain Machines IBM or your reseller specifies (called "Specific Machines"). International Business Machines Corporation or one of its subsidiaries ("IBM") owns copyrights in Code or has the right to license Code. IBM or a third party owns all copies of Code, including all copies made from them.

If you are the rightful possessor of a Specific Machine, IBM grants you a license to use the Code (or any replacement IBM provides) on, or in conjunction with, only the Specific Machine for which the Code is provided. IBM licenses the Code to only one rightful possessor at a time.

Under each license, IBM authorizes you to do only the following:

1. execute the Code to enable the Specific Machine to function according to its Official Published Specifications (called "Specifications");
2. make a backup or archival copy of the Code (unless IBM makes one available for your use), provided you reproduce the copyright notice and any other legend of ownership on the copy. You may use the copy only to replace the original, when necessary; and
3. execute and display the Code as necessary to maintain the Specific Machine.

You agree to acquire any replacement for, or additional copy of, Code directly from IBM in accordance with IBM's standard policies and practices. You also agree to use that Code under these terms.

You may transfer possession of the Code to another party only with the transfer of the Specific Machine. If you do so, you must 1) destroy all your copies of the Code that were not provided by IBM, 2) either give the other party all your IBM-provided copies of the Code or destroy them, and 3) notify the other party of these terms. IBM licenses the other party when it accepts these terms. These terms apply to all Code you acquire from any source.

Your license terminates when you no longer rightfully possess the Specific Machine.

1. Form Z125-4144

Actions you must not take

You agree to use the Code only as authorized above. You must not do, for example, any of the following:

1. Otherwise copy, display, transfer, adapt, modify, or distribute the Code (electronically or otherwise), except as IBM may authorize in the Specific Machine's Specifications or in writing to you;
2. Reverse assemble, reverse compile, or otherwise translate the Code unless expressly permitted by applicable law without the possibility of contractual waiver;
3. Sublicense or assign the license for the Code; or
4. Lease the Code or any copy of it.

Glossary

This glossary includes terms for the IBM FASTT Storage Manager Version 8.21 for UNIX and AIX environments.

This glossary also includes terms and definitions from: *IBM Dictionary of Computing*. New York: McGraw-Hill, 1994.

The following cross-reference convention is used in this glossary:

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

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

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

agent. A server program that receives virtual connections from the network manager (the client program) in an SNMP-TCP/IP network managing environment. The interface to a managed device.

auto volume transfer (AVT)/auto disk transfer (ADT). A built-in feature of the controller firmware that allows automatic logical drive-level failover rather than controller-level failover in case of controller failure on a storage subsystem. When AVT/ADT is enabled, the logical drive ownership will automatically be transferred to the alternate RAID controller if an error is detected on the path of the assigned preferred RAID controller of a logical volume. As soon as the path is re-established, the logical volumes are transferred back to their primary RAID controller. AVT/ADT works in conjunction with multipath I/O driver on the host system.

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

AWT. See *Abstract Windowing Toolkit*.

bootstrap protocol (BOOTP). A TCP/IP protocol used by a diskless workstation or network computer to obtain its IP address and other network information such as server address and default gateway. Upon startup, the client station sends out a BOOTP request in a UDP packet to the BOOTP server, which returns the required information. Unlike RARP, which uses only the layer 2 (Ethernet) frame for transport, the BOOTP request and response use an IP broadcast function that can send messages before a specific IP address is known.

BOOTP. See *bootstrap protocol*.

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.

dac. Disk array controller devices that represent the two controllers of the array.

dar. Disk array router that represents the entire array, including current and deferred paths to all LUNs (hdisks on AIX).

field replaceable unit. 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.

FRU. See *field replaceable unit*.

GUI. An acronym for graphical user interface.

hdisk. An AIX term; an hdisk represents a LUN on the array.

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

host computer. See *host*.

internet protocol (IP). In the Internet suite of protocols, a connection-less protocol that routes data through a network or interconnected networks and acts as an intermediary between the higher protocol layers and the physical network.

IP. See *internet protocol*.

JRE. An acronym for Java runtime environment.

LAN. See *local area network*.

logical unit number (LUN). An address that is used by the host to access a particular logical drive.

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

LUN. See *logical unit number*.

man pages. Manual pages. UNIX command. To get information about UNIX commands, type `man <command_name>`; for example, to get information about the `ls` command, type `man ls`.

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

MIB. See *management information base*.

NMS. An acronym for network management station.

NVSRAM. An acronym for nonvolatile storage random access memory.

ODM. An acronym for Object Data Manager.

PDF. An acronym for portable document format.

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

PTF. See *program temporary fix*.

RDAC. See *redundant dual active controller*.

redundant dual active controller (RDAC). Also known as redundant disk array controller, provides a multipath driver for a storage subsystem. Used with AIX and Solaris hosts.

router. A computer algorithm that determines the path of least resistance for traffic on a computer network.

SAM. System Administrator Manager.

SAN. See *storage area network*.

storage area network (SAN). Dedicated, high-speed networks of directly connected storage elements designed to move large amounts of data between host-independent distributed storage devices.

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

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

simple network management protocol (SNMP). A TCP/IP protocol that generally uses the user datagram protocol (UDP) to exchange messages between a management information base and a management client residing on a network. Because SNMP does not rely on the underlying communication protocols, it can be made available over other protocols, such as UDP/IP. 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 that are managed is defined and stored in the application's management information base (MIB).

SNMP. See *simple network management protocol and SNMPv1*.

SNMPv1. The original standard for SNMP is now referred to as SNMPv1.

subnet. A part of a network that is identified by a portion of an IP address.

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

TCP/IP. See *transmission control protocol/internet protocol*.

transmission control protocol/internet protocol (TCP/IP). A set of communication protocols that support peer-to-peer connectivity functions for both local and wide-area networks.

trap. A mechanism for SNMP agents to notify the SNMP management station of significant events.

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

WWN. See *worldwide name*.

Index

A

- about this document xi
- access volumes 3, 7
- actions you must not take 137
- address
 - for e-mail comments xiv
 - for reader comment form xiv
 - of IBM xiv
 - of the IBM director of licensing 135
- agent software
 - installing 38, 52
 - package 5
- agreement for licensed internal code 136
- AIX
 - configuration 15, 17
 - direct-attached configuration 3, 12
 - firmware
 - requirements 11
 - upgrading 16
 - verifying 14
 - hardware requirements 11, 123
 - host software
 - installing 13, 14
 - prerequisites 15
 - identifying
 - controller ID numbers 25
 - device names and bus numbers 26
 - logical drives by operating system device names 29
 - NVSRAM requirements 11
 - operating system requirements 9
 - prerequisites 14
 - RDAC driver
 - configuring 17
 - installing 15, 17
 - requirements 3, 14, 123
 - requirements
 - hardware 11, 123
 - host software 14
 - operating system 9
 - RDAC driver 3, 14, 123
 - software 123
 - system 123
 - restrictions 3, 5, 6, 7, 13
 - SMruntime installing 13
 - software requirements 123
 - system requirements 123
 - volumes, redistributing in case of failure 32
- alert notifications, configuring 62
- attributes
 - dac 20, 23
 - dar 18, 22
 - definitions 18
 - hdisk 20, 23, 24
 - LUN 20, 23, 24
 - RDAC driver 17
- audience of this document xi

- AVT/ADT, description 2

B

- books xii, 1
- bus numbers and device names, identifying 26, 44

C

- changing the host type 59
- client software
 - HP-UX
 - installing 38
 - requirements 36
 - package 3, 4
 - Solaris
 - installing 49
 - requirements 49
 - uninstalling 55
- cluster services
 - configurations 60
 - description 3
 - hardware requirements 67
 - HP-UX requirements 67
 - MC/Service Guard 67
 - Solaris requirements 67, 68
 - system dependencies 67
 - Veritas Cluster Server 67
 - Veritas Volume Manager 67, 68
- command line interface 2
- commands
 - fget_config 28
 - lsattr 28
 - lsdev 27
- comments about this document, how to send xiv
- configuration details, MC Service Guard 129
- configuration replication 2
- configuring
 - AIX 17
 - alert notifications 62
 - direct-attached 3, 12
 - RDAC driver 17, 68
 - Solaris with JNI host bus adapter cards 56
 - storage subsystems
 - AIX 3, 12, 15
 - HP-UX 39
 - procedures 60
 - Solaris 54
- controllers
 - diagnostics 2
 - firmware 1, 14, 40, 54
 - ID numbers, identifying 25, 44
 - IP addresses 9
 - properties 44
 - volumes 2
- creating a direct-attached configuration 12

D

- dac attributes 20, 23
- dar attributes 18, 22
- default partitioning, Solaris devices 55
- devices
 - identification 25
 - names and bus numbers, identifying 26, 44
 - Solaris, default partitioning for 55
- diagnostics, controller 2
- direct (out-of-band) management method 7
- direct-attached configuration 3, 12
- documentation
 - books xii
 - documents 1
 - files 1
 - online help xiii, 1, 62, 64
 - related publications xii
 - Web sites xii, xiii, 1
- drives, AIX, identifying by operating system device names 29

E

- e-mail address xiv, 63
- edition notice ii
- enabling multipath I/O with PV-links 41
- Enterprise Management window
 - description 4
 - online help 1
 - starting 60
- event monitoring 2

F

- failover protection 3
- failure support
 - cluster services 3, 67
 - RDAC driver 3, 5
 - redistributing volumes 32
- FAST, introduction to 9
- features of Storage Manager 8.21 2
- fget_config command 28
- filesets 123
- firmware
 - AIX
 - requirements 11
 - upgrading 16
 - verifying 14
 - description 1
 - HP-UX
 - requirements 35
 - upgrading 40
 - Solaris
 - requirements 47
 - upgrading 54
 - versions 40, 54
- FlashCopy 3, 77
 - command line interface 120
 - create drive 87
 - delete drive 111

FlashCopy (continued)

- disabling logical drive 106
- estimate repository life 84
- estimate repository logical drive capacity 83
- HP-UX instructions 116
- mapping drive to host 98
- missing logical drives 120
- parameters 79
- premium features 86
- progress of modification operation 105
- re-create logical drive 108
- re-size repository drive 109
- script editor 120
- script scenarios 77
- step-by-step guide 85
- UNIX regular disks 113
- view status 101

G

- GPFS 68

H

- HACMP 68
- hardware requirements
 - AIX 11, 123
 - HP-UX 35, 125
 - Solaris 47, 127
- hdisk
 - attributes 20
- hdisk attributes 20, 23, 24
- help xiii, 1, 62, 64
- heterogeneous environment 2, 60
- high-availability cluster services 3, 67, 68
- host software
 - AIX 14
 - AIX installing 13
 - packages 3
 - Solaris 50, 51
- host type 57, 59
- host-agent (in-band) management method 6
- how this document is organized xi
- how to send your comments xiv
- HP-UX
 - client software
 - installing 38
 - procedure 36, 37
 - requirements 36
 - cluster services, requirements 67
 - configuration of storage subsystem 39
 - firmware
 - requirements 35
 - upgrading 40
 - hardware requirements 35, 125
 - identifying
 - controller ID numbers 44
 - device names and bus numbers 44
 - NVSRAM
 - requirements 35
 - upgrading 40

HP-UX (*continued*)
operating system requirements 9
requirements
client software 36
cluster services 67
firmware 35
hardware 35, 125
NVS RAM 35
operating system 9
software 125
system 125
software requirements 125
system requirements 125
HyperTerminal version 9

I

IBM
address xiv
agreement for licensed internal code 136
director of licensing address 135
e-mail address xiv, 63
identifying
AIX
controller ID numbers 25
device names and bus numbers 26
logical drives by operating system device names 29
HP-UX
controller ID numbers 44
device names and bus numbers 44
in-band (host-agent) management method 6
installing
AIX
host software 13, 14
RDAC driver 15
SMruntime 13
host software 13, 36, 50
HP-UX
client software 38
host software 36
post-installation tasks 59
sequence of 3
SMruntime 13
Solaris
client software 49
host software 50
RDAC driver 51
utility software 52
IP addresses for FAS*T*T controllers 9

J

JNI
cards 56
settings 131

L

licensed internal code, agreement 136

logical drives
AIX, identifying by operating system device names 29
lsattr command 28
lsdev command 27
LUN attributes 23
LUNs
and access volumes 7
attributes 20, 23, 24

M

management method
direct (out-of-band) 7
host-agent (in-band) 6
management station 1
See storage management station
Mappings window 45
MC Service Guard 129
MC/Service Guard 67
missing logical drives 120
multipath I/O with PV-links 41

N

notices
edition ii
general 135
NVS RAM
AIX requirements 11
HP-UX
requirements 35
upgrading 40
Solaris
requirements 47
upgrading 54

O

online help xiii, 1, 62, 64
operating system
AIX 9, 123
HP-UX 9, 125
requirements
AIX 11
HP-UX 35
Solaris 47
Solaris 9, 127
organization of this document xi
out-of-band (direct) management method 7
overview of heterogeneous hosts 60

P

partitioning for Solaris devices 55
performing optional storage subsystem management tasks 64
prerequisites
AIX
firmware 11

prerequisites *(continued)*

AIX *(continued)*

- hardware 11, 123
- host software 14
- operating system 9
- RDAC driver 3, 14, 123
- software 123
- system 123

cluster services 67

HP-UX

- client software 36
- cluster services 67
- firmware 35
- hardware 35, 125
- NVSRAM 35
- operating system 9
- software 125
- system 125

Solaris

- client software 49
- cluster services 67, 68
- firmware 47
- hardware 47, 127
- host software 51
- NVSRAM 47
- operating system 9
- RDAC driver 3, 51
- software 127
- system 127
- utility software 51

products, developed 135

prohibited actions 137

PSSP 68

PV-links 41

R

RDAC driver

AIX

- configuring 17
- installing 15, 17
- requirements 3, 14, 123

attributes 17

description 3, 5

IDs 68

Solaris

- installing 51
- requirements 3, 51
- uninstalling 55

reader comment form xiv

README.txt file 1

redistributing volumes in case of failure 32

redundant dual active controller (RDAC) 3

related publications xii

renaming storage subsystem 64

requirements

AIX

- firmware 11
- hardware 11, 123
- host software 14
- operating system 9

requirements *(continued)*

AIX *(continued)*

- RDAC driver 3, 14, 123
- software 123
- system 123

cluster services 67

HP-UX

- client software 36
- cluster services 67
- firmware 35
- hardware 35, 125
- NVSRAM 35
- operating system 9
- software 125
- system 125

Solaris

- client software 49
- cluster services 67, 68
- firmware 47
- hardware 47, 127
- host software 51
- NVSRAM 47
- operating system 9
- RDAC driver 3, 51
- software 127
- system 127
- utility software 51

resources

- documents xii, 1
- e-mail address xiv, 63
- files 1
- online help xiii, 1, 62, 64
- Web sites xii, xiii, 1

restrictions

- AIX 3, 5, 6, 7, 13

S

sending your comments to IBM xiv

services offered in the U.S.A. 135

setting

- host type on Solaris 57
- IP addresses 9
- NMS for SNMP notification 63

SMagent

- agent software
- See agent software

SMclient

- client software
- See client software

SMruntime

- AIX installing 13
- HP-UX installing 36
- Solaris installing 48

SMutil

- utility software
- See utility software

SNMP notification 63

software package

- RDAC 5
- smagent 4

- software package *(continued)*
 - SMagent 5
 - SMutil 5
 - software requirements
 - AIX 123
 - HP-UX 125
 - Solaris 127
 - Solaris
 - client software
 - installing 49
 - procedure 49
 - requirements 49
 - uninstalling 55
 - cluster services requirements 67, 68
 - configuration of storage subsystem 54
 - devices, default partitioning 55
 - firmware
 - requirements 47
 - upgrading 54
 - hardware requirements 47, 127
 - host software
 - installing 50
 - procedure 51
 - requirements 51
 - JNI host bus adapter cards 56
 - NVSRAM
 - requirements 47
 - upgrading 54
 - operating system requirements 9
 - RDAC driver
 - installing 51
 - requirements 3, 51
 - uninstalling 55
 - requirements
 - client software 49
 - cluster services 67
 - firmware 47
 - hardware 47, 127
 - host software 51
 - NVSRAM 47
 - operating system 9
 - RDAC driver 3, 51
 - software 127
 - system 127
 - utility software 51
 - setting the host type 57
 - software requirements 127
 - system requirements 127
 - utility software
 - installing 52
 - requirements 51
 - uninstalling 55
 - starting the Enterprise Management window 60
 - starting the Subsystem Management window 63
 - storage management station 1
 - Storage Manager software
 - features 2
 - installation sequence 3
 - introduction 1
 - software packages 3
 - uninstalling, Solaris 55
 - storage subsystem
 - cluster services 67
 - configuration replication 2
 - configuring
 - AIX 15
 - HP-UX 39
 - procedures 60
 - Solaris 54
 - event monitoring 2
 - introduction 1
 - logical drives
 - identifying by operating system device names 29
 - management methods
 - direct (out-of-band) 7
 - host-agent (in-band) 6
 - management tasks 64
 - renaming 64
 - Subsystem Management window
 - description 5
 - figure of 63
 - online help 1
 - starting 63
 - system requirements
 - AIX 123
 - HP-UX 125
 - Solaris 127
- T**
- trademarks 135
- U**
- uninstalling Storage Manager software, Solaris 55
 - upgrading
 - AIX firmware 16
 - HP-UX
 - firmware 40
 - NVSRAM 40
 - Solaris
 - firmware 54
 - NVSRAM 54
 - utility software
 - package 5
 - Solaris
 - installing 52
 - requirements 51
 - uninstalling 55
- V**
- variable settings for JNI adapter cards 56
 - verifying AIX host firmware 14
 - Veritas Cluster Server 67
 - Veritas Volume Manager 67, 68
 - volumes
 - redistributing in case of failure 32

W

Web sites xii, xiii, 1

who should

 read this document xi

 use this document xi

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

IBM TotalStorage FAS*T* Storage Manager Version 8.21 for
UNIX and AIX Environments
Installation and Support Guide

Publication No. GC26-7490-01

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

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

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

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

Please tell us how we can improve this book:

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

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

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

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



Fold and Tape

Please do not staple

Fold and Tape



Printed in U.S.A.

GC26-7490-01



Spine information:



IBM TotalStorage FAS/T Storage
Manager Version 8.21 for
UNIX and AIX Environments

Installation and Support Guide