

IBM TotalStorage DS4000 Storage Manager 9



Installation and Support Guide for AIX, HP-UX, Solaris, and Linux on POWER

IBM TotalStorage DS4000 Storage Manager 9



Installation and Support Guide for AIX, HP-UX, Solaris, and Linux on POWER

Note:

Before using this information and the product it supports, read the information in “Notices” on page 169.

First Edition (January 2005)

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

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

Contents

Figures	vii
Tables	ix
About this document	xi
FAST product renaming	xi
Who should read this document	xii
DS4000 installation process overview	xii
DS4000 Storage Manager Version 9 library	xiv
How this document is organized	xiv
Notices used in this document	xv
Getting information, help and service	xv
Before you call	xv
Using the documentation	xvi
Web sites	xvi
How to send your comments	xvii
Chapter 1. Introduction	1
Product updates	2
Overview of DS4000 Storage Manager	3
The storage management station	3
DS4000 Storage Manager online help	4
What's new	4
Storage Manager 9.1 limitations	5
Overview of DS4000 Storage Manager software packages	5
SMruntime software package	6
SMclient software package	6
SMagent software package	7
SMutil software package	7
Failover drivers	7
Software installation sequence	7
Storage subsystem management methods	8
In-band (host-agent) management method	8
Out-of-band (direct) management method	10
Reviewing a sample network	12
Installation process diagram	13
Setting up IP addresses for DS4000 storage controllers	17
DS4300 storage server usage notes	18
DS4000 Storage Manager performance planning	18
Performance planning: Configuration examples	19
DS4000 Storage Manager premium features	19
Premium copy services features	20
Storage Partitioning premium feature	21
Enabling your premium features	21
Chapter 2. Installing and configuring storage management software on AIX host systems	23
Hardware, software and host system requirements	23
DS4000 hardware and firmware requirements	23
AIX system requirements	25
AIX restrictions	26
Installing DS4000 Storage Manager	27
Installing the client software	27

Installing SMagent (optional)	28
Configuring storage subsystems	29
Adding storage subsystems to SMClient	29
Upgrading controller firmware and NVSRAM	29
Storage Partitioning: Defining an AIX host group	33
Mapping LUNs to a storage partition	34
Creating direct and SAN-attached configurations	35
Creating a direct-attached configuration	35
Creating a SAN-attached configuration	35
Installing the RDAC failover driver	36
Prerequisites	36
Installing the RDAC driver	36
Identifying devices on AIX hosts	38
Performing initial device discovery	39
Verifying the installation and configuration	40
Identifying the controller ID numbers	40
Identifying device names and bus numbers	41
Identifying logical drives by operating system device names	46
Identifying FlashCopy logical drive types	47
Viewing and setting Object Data Manager (ODM) attributes	48
Attribute definitions	48
Using the lsattr command to view ODM attributes	54
Additional AIX configuration information	55
Disabling DS4000 cache mirroring	55
Using fast I/O failure for Fibre Channel devices	56
Using dynamic tracking of Fibre Channel devices	57
Using dynamic capacity expansion and dynamic volume expansion	57
Resolving disk array errors	59
Redistributing logical drives in case of failure	61
Chapter 3. Installing and configuring storage management software on	
HP-UX host systems	63
Hardware, firmware, and host system requirements	63
DS4000 hardware and firmware requirements	63
HP-UX system requirements	65
Installing DS4000 Storage Manager	66
Installing the client software	66
Installing the host software	68
Configuring storage subsystems	69
Adding storage subsystems to SMclient	70
Upgrading controller firmware and NVSRAM	71
Storage Partitioning: Defining an HP-UX host group	73
Mapping LUNs to a storage partition	74
Creating direct-attached and SAN-attached configurations	75
Creating a direct-attached configuration	75
Creating a SAN-attached configuration	75
Enabling multipath I/O with PV-links	75
Using PV-links: Method 1	76
Using PV-links: Method 2	78
Redistributing logical drives in case of failure	81
Chapter 4. Installing and configuring storage management software on	
Solaris host systems	83
Hardware, firmware, and host system requirements	83
DS4000 hardware and firmware requirements	83
Solaris system requirements	85

Installing DS4000 Storage Manager	86
Installing the client software	86
Installing the host software	88
Configuring storage subsystems	90
Adding storage subsystems to SMclient	91
Upgrading controller firmware and NVSRAM	91
Storage Partitioning: Defining a Solaris host group	93
Mapping LUNs to a storage partition	94
Creating direct-attached and SAN-attached configurations	95
Creating a direct-attached configuration	95
Creating a SAN-attached configuration	95
Configuring the host bus adapter cards	96
Installing the JNI adapter driver package	96
Modifying the JNI HBA settings	97
Installing the QLogic adapter driver package	102
Verifying the HBA installation	104
Modifying the QLogic HBA settings	104
Installing a failover driver on Solaris	105
Installing the RDAC failover driver	105
Installing the DMP failover driver	107
Verifying external storage	118
Redistributing logical drives in case of failure	118
Chapter 5. Using DS4000 with Linux on POWER-based hosts	121
Overview	121
Hardware, software and host system requirements	121
DS4000 hardware and firmware requirements	122
SUSE Enterprise Linux Server system requirements	124
Restrictions	124
Configuring storage subsystems on POWER-based Linux hosts	125
Adding storage subsystems to DS4000 SMClient	125
Upgrading DS4000 firmware and NVSRAM	126
Storage Partitioning: Defining a host group	128
Mapping LUNs to a DS4000 storage partition	131
Creating direct and SAN-attached configurations	132
Creating a direct-attached configuration	132
Creating a SAN-attached configuration	132
Configuring the host bus adapter drivers	133
Downloading the HBA driver and utility packages	133
Configuring HBA drivers	133
Performing initial device identification	134
Chapter 6. Completing the software configuration	139
Overview of heterogeneous host configurations	139
Configuring storage subsystems	139
Performing an initial automatic discovery of storage subsystems	139
Creating logical drives from free or unconfigured capacity	141
Configuring alert notifications	143
Renaming storage subsystems	144
Creating a storage subsystem profile	144
Performing other storage subsystem management tasks	144
Appendix A. MC/Service Guard configuration details	147
Appendix B. JNI and QLogic host bus adapter settings	149
JNI HBA card settings	149

QLogic HBA card settings	152
Appendix C. Downloading ESM firmware and drive firmware	153
Verifying and downloading ESM firmware.	153
Verifying current ESM firmware version	153
Downloading ESM firmware.	154
Downloading drive firmware.	156
Downloading Storage Manager version 9.10 drive firmware	156
Downloading Storage Manager version 8.x drive firmware	158
Appendix D. Additional DS4000 documentation	159
DS4500 Fibre Channel Storage Server library	159
DS4400 Fibre Channel Storage Server library	160
DS4300 Fibre Channel Storage Server library	161
DS4100 SATA Storage Server library	162
Other DS4000 and DS4000-related documents	163
Appendix E. Connecting HBAs in a FC switch environment	165
Appendix F. Accessibility	167
Notices	169
Trademarks.	170
Glossary	171
Index	181

Figures

1. Installation process flow by current publications	xiii
2. In-band managed storage subsystems	9
3. Out-of-band managed storage subsystems	11
4. Sample network using both out-of-band and in-band managed storage subsystems	13
5. Installation process (part 1)	14
6. Installation process (part 2)	15
7. Installation process (part 3)	16
8. DS4000/AIX system configuration.	38
9. Controller Properties window	41
10. Change Mappings window	47
11. Sample SMdevices command output (method 1)	77
12. Sample SMdevices command output (method 2)	78
13. Switch information report (part 1)	137
14. Switch information report (part 2)	138
15. Enterprise Management window.	140
16. Create Logical Drive — Introduction window	141
17. Fibre Channel fabric switch environment connecting host servers to a DS4000 Storage Server	165

Tables

1. Mapping of FASTT names to DS4000 Series names	xi
2. TotalStorage DS4000 Storage Manager Version 9 titles by user tasks	xiv
3. Current Storage Manager firmware versions by server and environment	1
4. Installation sequences of DS4000 Storage Manager software packages by host type	8
5. Maximum number of logical drives per partition.	18
6. Examples of possible configurations with more than 32 logical drives	19
7. Using Storage Partitioning with DS4000 storage servers	21
8. DS4000 storage servers that are used with AIX systems	24
9. Supported versions of hardware for AIX systems	24
10. Necessary criteria for concurrent download	30
11. File sets required for AIX 5.1 RDAC	37
12. File sets required for AIX 5.2 RDAC	37
13. File sets required for AIX 5.3 RDAC	37
14. Attributes for dar devices	49
15. Attributes for dac devices.	50
16. Attributes for hdisk devices	51
17. DS4000 storage servers that you can use with HP-UX systems	64
18. Supported versions of hardware for HP-UX systems	64
19. HP-UX kernel parameter configuration requirements	66
20. Sample record of logical drive preferred and alternate paths	79
21. DS4000 storage servers that you can use with Solaris systems.	84
22. Hardware you can use with Solaris systems	84
23. JNI model numbers, adapter drivers, and configuration files	96
24. Hardware that is supported with POWER-based SUSE Linux hosts	122
25. Hardware that is supported with POWER-based SUSE Linux hosts	123
26. Configuration settings for FCI-1063	149
27. Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562	149
28. Configuration settings for FCE-1063/FCE2-1063/FCE-6410/FCE2-6410	150
29. Configuration settings for FC64-1063	151
30. Configuration settings for QL2342	152
31. TotalStorage DS4500 Fibre Channel Storage Server document titles by user tasks	159
32. TotalStorage DS4400 Fibre Channel Storage Server document titles by user tasks	160
33. TotalStorage DS4300 Fibre Channel Storage Server document titles by user tasks	161
34. TotalStorage DS4100 SATA Storage Server document titles by user tasks	162
35. TotalStorage DS4000 and DS4000-related document titles by user tasks	163
36. DS4000 Storage Manager alternate keyboard operations	167

About this document

This document provides information about how to set up, install, configure, and work with the IBM® TotalStorage® DS4000 Storage Manager Version 9.1 in AIX®, HP-UX, Solaris, and Linux on POWER™.

Use this document to:

- Determine the hardware and software that is required to install the DS4000 Storage Manager software into your subsystem network
- Install the DS4000 Storage Manager software
- Upgrade the controller's nonvolatile storage random access memory (NVSRAM) and firmware
- Identify management features that are unique to your specific installation

Notes:

1. This document does not cover hardware installation or integration. For information about these topics, see the DS4000 Installation and Support Guide that is appropriate for your storage server. See Appendix D, "Additional DS4000 documentation," on page 159 for a list of these documents.
2. The controller firmware (version 06.10.xx.xx or higher) and NVSRAM that are included with the IBM DS4000 Storage Manager Version 9.1 support only the following DS4000 storage servers:
 - DS4300 standard (base) and Turbo options; Machine Type 1722; Models 60U/60X
 - DS4400; Machine Type 1742; Models 1RU/1RX
 - DS4500; Machine Type 1742; Models 90U/90X
3. You cannot use controller firmware version 6.10.xx.xx with the following DS4000 storage servers:
 - DS4100 — all models
 - DS4300 — SCU models
 - FAStT200
 - FAStT500

To find the firmware versions that you can use with these storage servers, see Table 3 on page 1. Please contact your IBM reseller or representative regarding any future availability of 6.10.xx.xx controller firmware for these storage servers.

FAStT product renaming

IBM is in the process of renaming some FAStT family products. Table 1 identifies each new DS4000 product name with its corresponding FAStT product name. Note that this change of **product name only** indicates no change in functionality or warranty. All products listed below with new names are functionally-equivalent and fully-interoperable. Each DS4000 product retains full IBM service as outlined in service contracts issued for analogous FAStT products.

Table 1. Mapping of FAStT names to DS4000 Series names

Current FAStT Product Name	New DS4000 Product Name
IBM TotalStorage FAStT Storage Server	IBM TotalStorage DS4000
FAStT	DS4000
FAStT Family	DS4000 Mid-range Disk System

Table 1. Mapping of FAStT names to DS4000 Series names (continued)

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

Who should read this document

This document is intended for system administrators and storage administrators who are responsible for installing storage management software.

Readers should have knowledge of the following concepts:

- Redundant array of independent disks (RAID)
- Small computer system interface (SCSI)
- Fibre Channel and serial ATA (SATA) technology

Readers should also have working knowledge of the applicable host operating systems that are used with DS4000 Storage Manager.

DS4000 installation process overview

Figure 1 on page xiii gives an overview of the DS4000 hardware and the DS4000 Storage Manager software installation process. Lined arrows in the flow chart indicate consecutive hardware and software installation process steps. Labeled arrows indicate which document provides detailed information about a step.

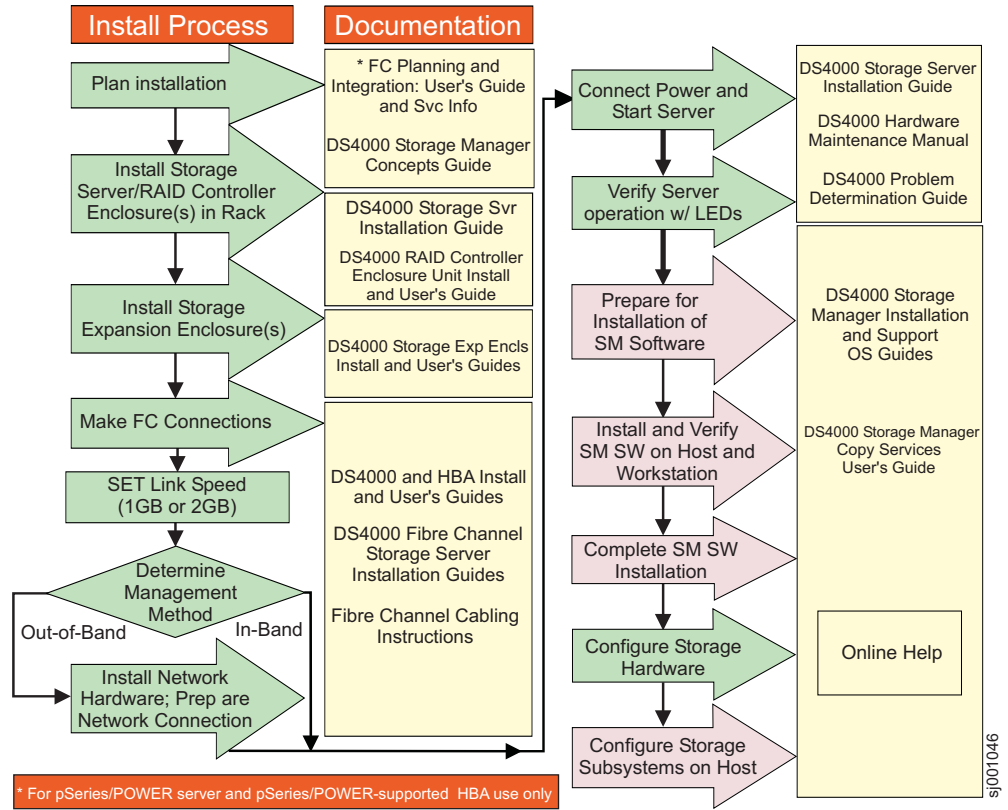


Figure 1. Installation process flow by current publications

DS4000 Storage Manager Version 9 library

Table 2 presents an overview of DS4000 Storage Manager product library documents by the common tasks that they address.

You can access the documents listed in this table at one of the following Web sites:

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

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

Table 2. TotalStorage DS4000 Storage Manager Version 9 titles by user tasks

Title	User tasks					
	Planning	Hardware installation	Software installation	Configuration	Operation and administration	Diagnosis and maintenance
<i>IBM TotalStorage DS4000 Storage Manager Version 9 Installation and Support Guide for Windows 2000/Server 2003, NetWare, ESX Server, and Linux, GC26-7706</i>	✓		✓	✓		
<i>IBM TotalStorage DS4000 Storage Manager Version 9 Installation and Support Guide for AIX, UNIX, Solaris and Linux on POWER, GC26-7705</i>	✓		✓	✓		
<i>IBM TotalStorage DS4000 Storage Manager Version 9 Copy Services User's Guide, GC26-7707</i>	✓		✓	✓	✓	
<i>IBM TotalStorage DS4000 Storage Manager Version 9 Concepts Guide, GC26-7734</i>	✓	✓	✓	✓	✓	✓

How this document is organized

Chapter 1, "Introduction," on page 1 provides an introduction to the DS4000 Storage Manager product, including information about product resources in addition to this document.

Chapter 2, "Installing and configuring storage management software on AIX host systems," on page 23 provides information about how to install and configure DS4000 Storage Manager with an AIX host operating system.

Chapter 3, “Installing and configuring storage management software on HP-UX host systems,” on page 63 provides information about how to install and configure DS4000 Storage Manager with an HP-UX host operating system.

Chapter 4, “Installing and configuring storage management software on Solaris host systems,” on page 83 provides information about how to install and configure DS4000 Storage Manager with a Solaris host operating system.

Chapter 5, “Using DS4000 with Linux on POWER-based hosts,” on page 121 provides information about how to configure a DS4000 subsystem with a SUSE Linux Enterprise Server operating system on a POWER or pSeries-based host.

Chapter 6, “Completing the software configuration,” on page 139 provides step-by-step instructions for post-installation tasks, such as how to create storage arrays, logical drives, and storage partitions.

Appendix A, “MC/Service Guard configuration details,” on page 147 provides the procedure for how to correct the primary and alternate paths of the imported volume groups that are changed after using `vgimport -m -s` with LVM commands.

Appendix B, “JNI and QLogic host bus adapter settings,” on page 149 provides the correct JNI and QLogic host bus adapter (HBA) settings for Solaris host operating systems.

Appendix C, “Downloading ESM firmware and drive firmware,” on page 153 provides the procedures for verifying and downloading Environmental Service Module (ESM) firmware.

Appendix E, “Connecting HBAs in a FC switch environment,” on page 165 provides an example of a Fibre Channel fabric switch configuration, and a Web site that contains links to switch documentation.

Notices used in this document

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

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

Getting information, help and service

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

Before you call

Before you call, take these steps to try to solve the problem yourself:

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

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

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the DS4000 Storage Manager online help or in the documents that are provided with your system and software. The information that comes with your system also describes the diagnostic tests that you can perform. Most servers, operating systems, and programs come with information that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the information for the operating system or program.

Using the documentation

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

Web sites

The most up-to-date information about DS4000 storage servers and DS4000 Storage Manager, including documentation and the most recent software, firmware, and NVSRAM downloads, can be found at the following Web sites.

DS4000 Fibre Channel storage servers

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

<http://www-1.ibm.com/servers/storage/disk/index.html>

IBM TotalStorage products

Find information about all IBM TotalStorage products:

www.storage.ibm.com/

TotalStorage DS4000 interoperability matrix

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

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

DS4000 Storage Manager readme files

Find the latest readme files for DS4000 Storage Manager:

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

Click the link for your DS4000 storage server. When the page opens, click the Download tab. Click the link for **Current recommended Firmware and Storage Manager**. In the tables, find the Storage Manager listing for your operating system and click the **v9.1x** link in the Current Version column.

Fix delivery center for AIX and Linux on POWER

Find the latest AIX and Linux on POWER information and downloads:

www-912.ibm.com/eserver/support/fixes/fcgui.jsp

In the **Server** drop down menu, select **pSeries family**. Then select your product or fix type.

Linux on pSeries support

Find information about using Linux with pSeries:

www.ibm.com/servers/eserver/pseries/linux/

Linux on POWER resource center

Find information about using Linux with POWER:

www.ibm.com/servers/enable/linux/power/

Storage Area Network (SAN) support

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

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

DS4000 technical support

Find DS4000 downloads, hints and tips, documentation, parts information, HBA and Fibre Channel support:

<http://www-1.ibm.com/servers/storage/support/disk/>

Premium feature activation

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

[www.storage.ibm.com/pfeatures.html/](http://www.storage.ibm.com/pfeatures.html)

IBM publications center

Find IBM publications:

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

How to send your comments

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

- E-mail

Submit your comments electronically to:

starpubs@us.ibm.com

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

- Mail

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

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

Chapter 1. Introduction

IBM DS4000 Storage Manager version 9.1x host software supports new features that are incorporated into the latest release of controller firmware version 06.10.xx.xx. IBM DS4000 Storage Manager version 9.1x also manages DS4000 storage servers with all the generally-released controller firmware versions, from 04.01.xx.xx through 05.4x.xx.xx.

See Table 3 for information on latest supported controller firmware code levels available for all DS4000 storage server models.

All of the controller firmware versions that are listed in Table 3 are available free-of-charge. However, you must purchase an EXP100 storage expansion enclosure or DS4100 storage server to obtain controller firmware code versions 05.41.xx.xx or 05.42.xx.xx, respectively.

Note: New features that are associated with the 06.10.xx.xx controller firmware are listed in “What’s new” on page 4.

Table 3. Current Storage Manager firmware versions by server and environment

Storage server name	Machine type	Model number	Latest controller firmware version for Fibre Channel environment	Latest controller firmware version for SATA environment
IBM Netfinity Fibre Channel RAID Controller Unit	3526	1RU 1RX	04.01.xx.xx	N/A
DS4100	1724	100	N/A	05.42.xx.xx
FAStT200 High Availability (HA)	3542	2RU 2RX	05.30.xx.xx	N/A
FAStT200		1RU 1RX		
FAStT500	3552	1RU 1RX		
DS4300 Single Controller Unit (SCU)	1722	6LU 6LX	05.34.xx.xx	05.41.xx.xx
DS4300 Base		60X 60U	06.1x.xx.xx	06.1x.xx.xx
DS4300 Turbo				
DS4400	1742	1RU 1RX		06.1x.xx.xx
DS4500		90U 90X		06.1x.xx.xx

This document provides you with information about how to install, configure, and work with DS4000 Storage Manager in the following host environments:

- AIX
- HP-UX
- Solaris
- SUSE Linux Enterprise Server on POWER

Note: For information about installing Storage Manager on other host operating systems, see the *IBM TotalStorage DS4000 Storage Manager 9 Installation and Support Guide for Windows 2000/Server 2003, NetWare, ESX Server, and Linux*.

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

readme files

Read these first.

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

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

Click the link for your storage server.

When the storage server page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the resulting download files matrix, the links to the firmware downloads also provide the readme files for Storage Manager or for your host operating system.

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

IBM DS4000 Storage Manager Concepts Guide

Use this reference document to become familiar with the terminology and the features of the DS4000 Storage Manager software. This document is available on the DS4000 Storage Manager installation CD and at the Web site listed above.

Product updates

Important

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

Download the latest version of the DS4000 Storage Manager host software and the any appropriate DS4000 product series firmware at the time of the initial installation and when product updates become available.

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

www.ibm.com/support/us/

Go to the **Personalized Support** section of the web page and click **My support**. On the next page, if you have not already done so, register to use the site by clicking **register now**.

Perform the following steps to receive product updates:

1. After you have registered, type your user ID and password to log into the site. The **My support** page opens.

2. Click **add products**. Using the pull downs in the **Products** area, select **Storage** → **Computer Storage** → **Disk Storage Systems** → **TotalStorage DS4000 Midrange Disk Systems & FAST Stor Svrs**.
3. Place a check in the box for the machine type of your DS4000 series product, as well as any other attached DS4000 series product(s) for which you would like to receive information. Select **Add products**. The **My support** page reopens.
4. Select **Subscribe to email**. In the **Documents** area on the next page, use the pull down and select **Storage**.
5. On the next page, place a check in the following boxes:
 - a. **Please send these documents by weekly email**
 - b. **Downloads and drivers**
 - c. **Flashes**and any others you may be interested in, and then click **Update**.

Overview of DS4000 Storage Manager

IBM TotalStorage DS4000 Storage Manager consists of a set of client and host tools that enable you to manage IBM DS4300 standard (base) and Turbo options, DS4400, and DS4500 Fibre Channel Storage Server subsystems from a storage management station.

Note: DS4300 single-controller unit (SCU) storage servers are not supported with controller firmware version 6.1x.xx.xx. However, you can use DS4300 standard (base) and Turbo storage server options with controller firmware 6.1x.xx.xx.

DS4000 Storage Manager software is available on the product CD. You can also download DS4000 Storage Manager software from the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level.

The storage management station

The *storage management station* is the system that is responsible for managing all, or a portion of, a storage network. It communicates with the network management agents that reside in the managed nodes using a network management protocol, such as Simple Network Management Protocol (SNMP).

Storage management commands are sent to the storage subsystem controllers, where the controller firmware validates and runs the commands, and then returns status and configuration information to the client software.

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

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

Note: Even though you can install the storage management software on a host, the host still uses the Transmission Control Protocol/Internet Protocol (TCP/IP) to communicate with the host-agent. The host-agent communicates with the controller over the Fibre Channel connection through the *access volume*.

DS4000 Storage Manager online help

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

You can access the help systems from the Enterprise Management and Subsystem Management windows in DS4000 Storage Manager 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.

What's new

This section describes new DS4000 Storage Manager features and new host operating system capabilities that are available with Storage Manager.

FAST product renaming

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

DS4000 Storage Manager 9.1 provides the following new features:

- **Enhancements to the Remote Mirror Option premium feature**

With the introduction of the asynchronous write mode feature in DS4000 Storage Manager 9.1, the Remote Mirror Option is now referred to as the Enhanced Remote Mirror Option..

See *IBM TotalStorage DS4000 Storage Manager 9 Copy Services Guide* for details about changes to the Enhanced Remote Mirror Option.

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

- **Staged controller firmware download**

You can download controller firmware from Storage Manager to all DS4000 controllers for later activation. Depending on your firmware version, DS4000 storage server model, and host operating system, the following options might be available:

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

Note: Staged controller firmware download is not supported on DS4400 storage servers.

- **Parallel drive firmware download**

Hard drive firmware packages can be downloaded to multiple drives simultaneously, which minimizes downtime. In addition, all files that are associated with a firmware update are now bundled into a single firmware package.

See the Subsystem Management window online help for drive firmware download procedures.

Note: This is not the same thing as concurrent download, which is not supported for drive firmware. You must stop I/O before downloading drive firmware.

- **Subsystem Management Window menu enhancements**

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

- Maintenance
- Troubleshooting
- Recovery

- **DS4000 support for POWER-based Linux hosts**

You can use POWER-based Linux hosts in a DS4000 subsystem. However, the DS4000 Storage Manager client software is not currently available for POWER-based Linux hosts.

To manage DS4000 storage subsystems with your POWER-based Linux host, you must install the Storage Manager client software (SMclient) on an AIX 5L, Solaris, HP-UX, Windows or i386 processor-based Linux server, and use that server as your storage management station.

- **Full command-line interface capability**

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

Storage Manager 9.1 limitations

Storage Manager 9.1-level controller firmware version 06.1x.xx.xx and corresponding NVSRAM are not supported on the following storage servers.

- DS4100 (1724-100)
- FAStT200 (3542, all models)
- FAStT500 (3552, all models)
- DS4300 (1722-6LU/6LX [single-controller unit only])

Therefore, the Storage Manager features that are specific to version 06.1x.xx.xx firmware are unavailable for these DS4000 servers. However, IBM recommends that you do use Storage Manager 9.1 host software to manage these DS4000 servers, using controller firmware versions 04.10.xx.xx or higher. For more information about which firmware to use with your configuration, see Table 3 on page 1.

Overview of DS4000 Storage Manager software packages

DS4000 Storage Manager contains the following client software packages:

SMruntime software

DS4000 Storage Manager Java™ compiler

SMclient software

DS4000 Storage Manager client package

DS4000 Storage Manager contains the following host software packages:

SMagent software

DS4000 Storage Manager agent package

SMutil software

DS4000 Storage Manager utility package

RDAC

DS4000 Storage Manager multipath device drivers

Note: In addition to this document, see the *IBM DS4000 Storage Manager Concepts Guide* for more information about DS4000 Storage Manager software.

SMruntime software package

The DS4000 Storage Manager runtime software, SMruntime, is a Java compiler for the DS4000 Storage Manager client software, SMclient. SMruntime must be installed before SMclient is installed.

SMclient software package

The DS4000 Storage Manager client software, SMclient, is a Java-based graphical user interface (GUI).

SMclient enables you to configure, manage and monitor DS4000 storage servers and storage expansion enclosures in a storage subsystem, either through a host system or through a storage management station.

Specifically, SMclient enables you to perform the following tasks:

- Configure disk arrays and logical volumes
- Assign names to arrays and volume groups
- Assign logical volumes to storage partitions
- Replace and rebuild failed disk drives
- Expand the size of arrays
- Expand the size of logical volumes
- Change RAID-type arrays
- Configure and add additional host partitions with the premium partitioning feature
- Monitor the status of DS4000 storage servers
- Perform troubleshooting and maintenance tasks, such as downloading controller, ESM and drive firmware and NVSRAM
- View the major events log (MEL)
- Assign redundant RAID controllers
- Expand storage capacity using any of the following storage expansion enclosures:
 - DS4000 EXP500 Fibre Channel storage expansion enclosure
 - DS4000 EXP710 Fibre Channel storage expansion enclosure
 - DS4000 EXP700 Fibre Channel storage expansion enclosure
 - DS4000 EXP100 SATA storage expansion enclosure

Note: For important information about using the DS4000 EXP100 SATA storage expansion enclosure, see *IBM TotalStorage DS4000 EXP100 storage expansion enclosure Installation and User's Guide* (GC26-7601).

The SMclient contains two main components:

Enterprise Management

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

Subsystem Management

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

SMagent software package

The DS4000 Storage Manager software agent package, SMagent, contains optional host-agent software, which you can use on AIX, HP-UX, and Solaris host systems to manage storage subsystems through the host Fibre Channel connection.

Note: SMagent is not supported with controller firmware version 5.42.xx.xx.

SMagent takes requests from a storage management station that is connected to the host through a network connection, and passes the requests to the storage subsystem controllers through the Fibre Channel I/O path.

For more information about managing storage subsystems through SMagent, see “In-band (host-agent) management method” on page 8.

SMutil software package

You can use the DS4000 Storage Manager utility, SMutil, to register and map new logical drives to the operating system.

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

Failover drivers

Solaris host systems requires one of the following failover drivers for Fibre Channel path redundancy:

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

The failover driver monitors I/O paths. If a component failure occurs in one of the Fibre Channel paths, the failover driver reroutes all I/O to another path.

Note: The AIX RDAC driver files are not included on the DS4000 Storage Manager installation CD. To install them, you must download the files from the Web site listed in “Installing the RDAC failover driver” on page 36 and follow the instructions in that section.

Software installation sequence

Install the DS4000 Storage Manager software packages in the sequences shown in Table 4 on page 8.

Table 4. Installation sequences of DS4000 Storage Manager software packages by host type

Step	AIX	HP-UX	Solaris
1	SMruntime	SMruntime	SMruntime
2	SMclient	SMclient	SMclient
3	SMagent	SMagent	SMagent
4	RDAC	SMutil	SMutil
5			RDAC or DMP
Location of procedures	“Installing DS4000 Storage Manager” on page 27	“Installing DS4000 Storage Manager” on page 66	“Installing DS4000 Storage Manager” on page 86

Note: DS4000 Storage Manager client software is not currently available for SUSE Linux Enterprise Server on POWER-based hosts. If you are using a POWER-based SUSE Linux host with your DS4000 subsystem, you must install the Storage Manager client software on an AIX 5L, HP-UX, Solaris, Windows or i386 processor-based Linux server, and use that server as your storage management station.

:

Storage subsystem management methods

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

In-band (host-agent) management method

Using this method, you manage the storage subsystems through the Fibre Channel I/O path to the host.

Out-of-band (direct) management method

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

The management methods are described in the next sections.

In-band (host-agent) management method

When you use the in-band management method, you manage the storage subsystems through the Fibre Channel I/O path to the host.

The management information can be processed by the host or passed to the storage management station through the network connection. Figure 2 on page 9 shows the in-band management method.

Restrictions:

- If both Fibre Channel connections to the controllers are lost, the SMclient software cannot be accessed for problem determination.
- In-band management is not supported with controller firmware version 5.42.xx.xx.




Figure 2. In-band managed storage subsystems

In-band management method advantages

Using the in-band management method provides the following advantages:

- You do not need to run Ethernet cables to the controllers.
- You do not need a Dynamic Host Configuration Protocol (DHCP) bootstrap protocol (BOOTP) server to connect the storage subsystems to the network.
- You do not need to configure the controller network (described in Chapter 2, “Installing and configuring storage management software on AIX host systems,” on page 23, Chapter 3, “Installing and configuring storage management software on HP-UX host systems,” on page 63, or Chapter 4, “Installing and configuring storage management software on Solaris host systems,” on page 83).
- 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.

In-band management method disadvantages

Using the in-band management method has the following disadvantages:

- If both Fibre Channel connections to the controllers are lost, the SMclient software cannot be accessed for problem determination.

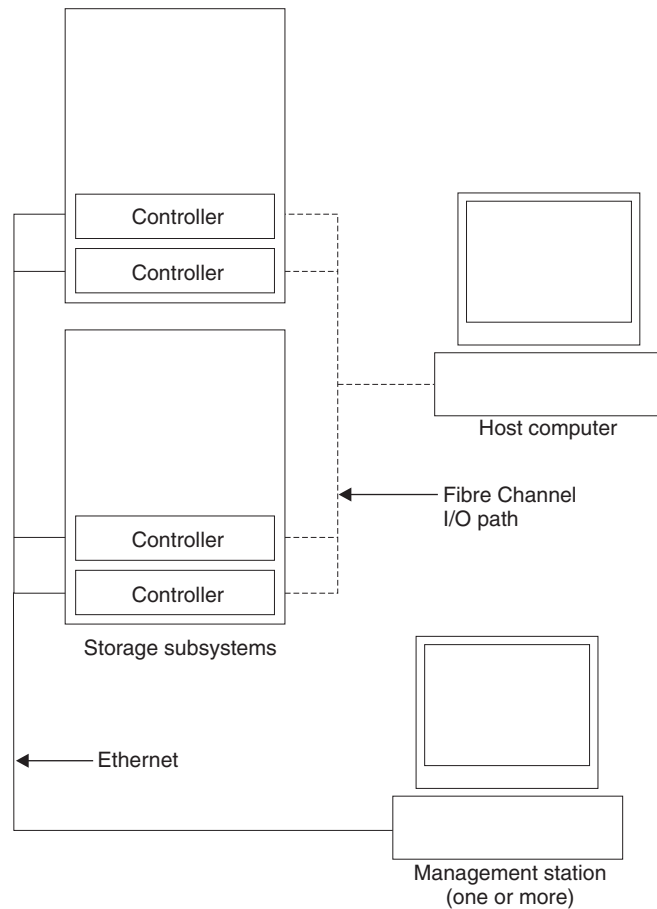
- Unlike out-of-band management, in which the Fibre Channel connections are dedicated for I/O, in-band management uses the Fibre Channel for both I/O and status reporting. Therefore, in-band management uses more bandwidth on the Fibre Channel than out-of-band management, and this can have an impact on performance.
- You are limited to configuring one less LUN than the maximum number allowed by the operating system and host adapter that you are using.
- The host-agent requires a special logical drive, called an *access volume*, to communicate with the controllers in the storage subsystem.

Important: The access volume uses one of the LUNs. If your host already has the maximum number of LUNs configured, either use the out-of-band 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.

Out-of-band (direct) management method

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

To manage the storage subsystem through the Ethernet connections, you must define the IP address and host name for each controller and attach a cable to the Ethernet ports on each of the storage subsystem controllers. Figure 3 on page 11 shows the out-of-band management method.



SJ000708

Figure 3. Out-of-band managed storage subsystems

If you use the out-of-band management method, see “Setting up IP addresses for DS4000 storage controllers” on page 17.

Out-of-band management method advantages

Using the out-of-band management method provides the following advantages:

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

Out-of-band management method disadvantages

Using the 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.

Reviewing a sample network

Figure 4 on page 13 shows an example of a network that contains both an out-of-band storage subsystem (Network A) and an in-band managed storage subsystem (Network B).

Out-of-band-managed storage subsystem: Network A is an out-of-band-managed storage subsystem. Network A contains the following components:

- A DHCP/BOOTP server
- A network management station (NMS) for Simple Network Management Protocol (SNMP) traps
- A host that is connected to a storage subsystem through a fibre-channel I/O path
- A management station that is connected by an Ethernet cable to the storage subsystem controllers

Note: If the controllers, static IP addresses, or default IP addresses are used, you do not need to setup the DHCP/BOOTP server.

In-band storage subsystem: Network B is an in-band-managed storage subsystem. Network B contains the following components:

- A host that is connected to a storage subsystem through a fibre-channel I/O path
- A management station that is connected by an Ethernet cable to the host computer

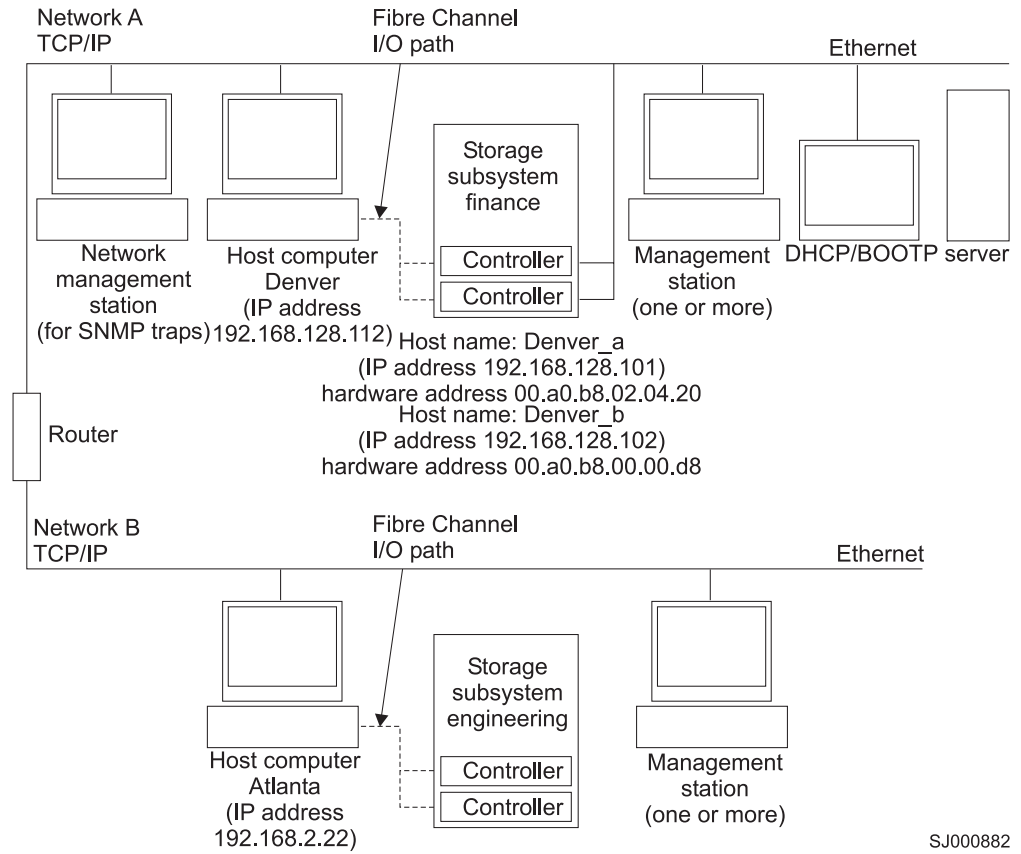
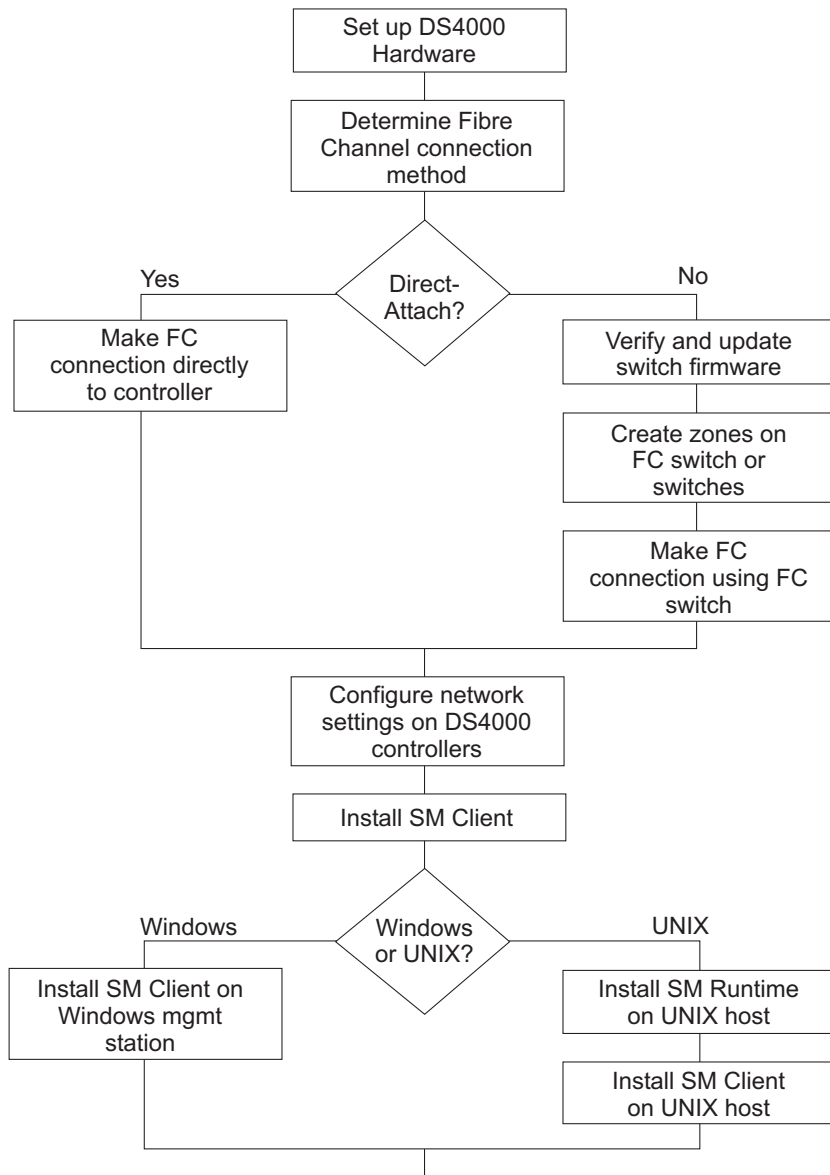


Figure 4. Sample network using both out-of-band and in-band managed storage subsystems

Installation process diagram

Figure 5 on page 14, Figure 6 on page 15, and Figure 7 on page 16 provide an overview of the installation process for out-of-band management.

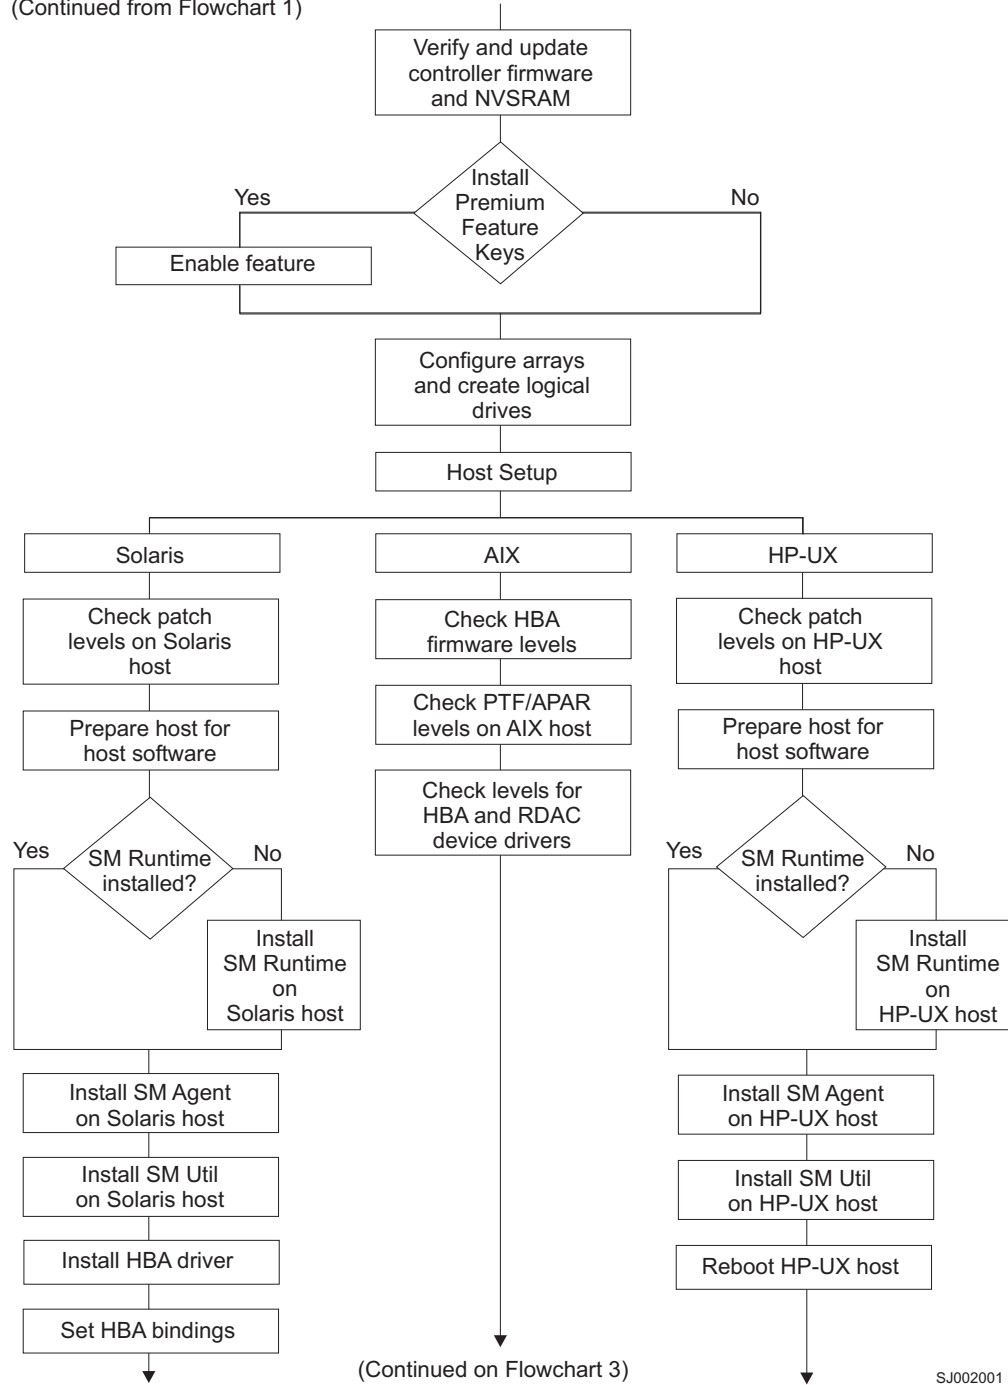


(Continued on Flowchart 2)

SJ002000

Figure 5. Installation process (part 1)

(Continued from Flowchart 1)




(Continued on Flowchart 3)

SJ002001

Figure 6. Installation process (part 2)

(Continued from Flowchart 2)



SJ002002

Figure 7. Installation process (part 3)

Setting up IP addresses for DS4000 storage controllers

To use the out-of-band management method without setting up a DHCP BOOTP server, you must assign IP addresses to the DS4000 controllers using command-line interface (CLI) commands through serial cables that are connected to a terminal emulator.

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

Ensure that you have the following addresses:

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

Complete the following steps to set up the DS4000 controller IP addresses by using serial ports:

1. Stop all I/O to the DS4000 controllers.
2. Connect a null modem serial cable from one of the controllers to a system with a terminal emulator available.
3. Open the HyperTerminal and from the menu bar, click **File** → **Properties** → **Configure**. Choose the following settings:

Bits/second 57600

Attention: Choosing a different setting causes undesirable results.

Data bits 8

Parity None

Stop Bit 1

Flow Control XON/XOFF

4. Connect to the DS4000 storage server and send a break signal (Ctrl+Break for most emulators).
5. Repeat this step until the following message is displayed:

Press the space bar for baud rate within 5 seconds.

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

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

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

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

- Controller A = 192.168.10.101

- Controller B = 192.168.10.102
 - IP address mask = 255.255.255.0
11. Type netCfgSet to change the network configuration information.

Note: Press Enter to advance to the next field. Type the new IP address in the **My IP Address** field.
 12. Assign an IP address to the controller.
 13. Type netCfgShow to verify the new network settings.
 14. Disconnect from the first controller and connect to the second controller.
 15. Repeat step 1 on page 17 through step 14 to assign the second IP address to the second controller.
 16. Turn the controller unit off and on to restart the storage server.

DS4300 storage server usage notes

Note the following restrictions:

- DS4300 storage servers do not support DS4000 EXP500 Fibre Channel storage expansion enclosures.
- DS4300 single-controller (SCU) storage servers are not supported with controller firmware 6.1x.xx.xx. However, you can use DS4300 standard (base) and Turbo storage servers with controller firmware 6.1x.xx.xx.

DS4000 Storage Manager performance planning

DS4000 storage systems provide a high degree of flexibility to meet today's demanding storage applications. As such, care should be taken to plan configurations that meet specific application performance requirements.

Like most storage vendors, IBM publishes data sheets and specifications for the capabilities of the storage subsystem. These are intended for use as reference points for raw capabilities, but do not necessarily imply that actual performance specifications will be achieved for all configurations and operating environments. Many factors that are external to the DS4000 storage server, such as server processor and memory, host bus adapters, command queue depth settings, and I/O size, can impact performance.

You must plan the configuration and setup of the DS4000 for your operating environment. If proper planning is not performed, poor performance can result, such as low throughput or high disk latencies (response times).

If your system uses controller firmware versions 5.4.xx.xx or later, you can create storage partitions with greater than 32 logical drives (up to 256). Take this factor into account during configuration planning. Depending on the I/O load requirements, it might be beneficial to spread volumes out on multiple partitions, if your operating system allows it.

Note: Not all operating system environments can support 256 logical drives per partition. See Table 5:

Table 5. Maximum number of logical drives per partition

Operating system	Maximum number of LUNs per partition
AIX	256
HP-UX 11.0	32

Table 5. Maximum number of logical drives per partition (continued)

Operating system	Maximum number of LUNs per partition
HP-UX 11.i	128
Solaris	256 with DMP 32 with RDAC
SUSE Linux Enterprise Server 9	256

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

Performance planning: Configuration examples

If your system is running DS4000 Storage Manager firmware version 5.4.xx.xx or later, you can configure all 64 logical disks into one storage partition with two HBAs. You can also configure 256 logical disks into one partition with two HBAs. See Table 6:

Note: Systems running DS4000 Storage Manager firmware version 5.3.xx.xx or earlier require two storage partitions, each with 32 logical drives and four HBAs, to connect 64 logical disks to a host.

Table 6. Examples of possible configurations with more than 32 logical drives

Configuration number	Storage Manager firmware version	HBAs	Storage partitions	Logical drives
1	5.3.xx.xx	4	2	64
2	5.4.xx.xx	2	1	64
3	5.4.xx.xx or 6.10.xx.xx	2	1	256

You might expect the following performance results with the configurations listed in Table 6:

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

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

- *IBM TotalStorage DS4000 Storage Manager Concepts Guide*
- *IBM TotalStorage: DS4000 Best Practices Guide*

DS4000 Storage Manager premium features

DS4000 Storage Manager supports the following premium features, which are available for purchase separately from IBM or an IBM Business Partner:

Premium copy services features

The following copy services are available with Storage Manager 9.1:

- FlashCopy®
- VolumeCopy
- Enhanced Remote Mirror Option

Storage Partitioning premium feature

Storage Partitioning is standard on all DS4000 storage servers that are supported by DS4000 Storage Manager controller firmware version 6.10.xx.xx.

Premium copy services features

This document lists the premium copy services features that are supported by DS4000 Storage Manager, but does not describe them in detail. For detailed copy services concepts and procedures, see the following documents, which are included on the DS4000 Storage Manager installation CD-ROM.

- *IBM TotalStorage DS4000 Storage Manager Copy Services User's Guide*
- *IBM TotalStorage DS4000 Storage Manager Concepts Guide*

DS4000 Storage Manager supports the following premium copy services.

FlashCopy

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

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

VolumeCopy

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

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

Notes:

1. You must have FlashCopy enabled in order to use VolumeCopy. VolumeCopy can be purchased as part of a package with FlashCopy, or it can be purchased at a later time as an enhancement to FlashCopy.
2. VolumeCopy is not supported with Storage Manager controller firmware version 5.42.xx.xx, or with the DS4100 storage server.

Enhanced Remote Mirror Option

The Enhanced Remote Mirror Option provides online, real-time replication of data between storage subsystems over a remote distance.

In the event of a disaster or unrecoverable error at one storage subsystem, the Enhanced Remote Mirror Option enables you to promote a second storage subsystem to take over responsibility for normal I/O operations.

Notes:

1. With Storage Manager controller firmware versions earlier than 6.10.xx.xx, you cannot use the Enhanced Remote Mirror Option.

However, you can use a previous version of the Remote Mirror Option. See *IBM TotalStorage DS4000 Storage Manager Copy Services User's Guide* for more information about the different versions.

2. **Exception:** You cannot use the Remote Mirror Option or the Enhanced Remote Mirror Option with Storage Manager controller firmware versions 5.41.xx.xx or 5.42.xx.xx, or with the DS4100 storage server.

Storage Partitioning premium feature

The Storage Partitioning feature enables you to associate a set of logical drives on a DS4000 storage server that can only be accessed by specified hosts and host ports. This association of logical drives to a set of hosts and host ports is called a *storage partition*. The benefit of defining storage partitions is to allow controlled access to the logical drives to only those hosts also defined in the storage partition.

Table 7 shows how many storage partitions are standard on your DS4000 storage server.

Table 7. Using Storage Partitioning with DS4000 storage servers

DS4000 Storage Server	Storage Partitioning enabled by default	Maximum number of storage partitions
DS4100	No	16
DS4500	Yes (16 partitions standard)	64
DS4400	Yes (64 partitions standard)	64
DS4300 with Turbo option	Yes (8 partitions standard)	64

If you do not use Storage Partitioning, all logical drives are contained in a default host group. The logical drives in the default host group can be accessed by any Fibre Channel initiator that has access to the DS4000 host port.

When the DS4000 storage server is attached to a SAN, you can use zoning within the fabric to limit access to the DS4000 host ports to specific set of hosts. Also, when homogeneous host servers are directly attached to the DS4000 storage server, access to all logical drives might be satisfactory.

Note: For more information about fabric switches, see Appendix E, “Connecting HBAs in a FC switch environment,” on page 165.

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

Enabling your premium features

Complete the following steps to enable a premium feature:

1. Gather the following information:
 - The feature activation code that accompanies the premium feature
 - Your controller unit machine type, model number and IBM serial number, which are printed on the label on the back of your DS4000 controller unit
 - Your 32-digit DS4000 Subsystem feature enable identifier, which you can view in the DS4000 Storage Manager Subsystem Management Window by clicking **Storage Subsystem** → **Premium Features** → **List**

2. Ensure that your controller unit and storage expansion enclosures are connected, powered on, and configured.
3. Generate the feature key file and enable your premium feature by completing the steps at the following Web site:

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

Your system is now ready to use the premium feature.

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

Chapter 2. Installing and configuring storage management software on AIX host systems

This chapter provides the following information for AIX operating systems:

- “Hardware, software and host system requirements”
- “AIX restrictions” on page 26
- “Installing DS4000 Storage Manager” on page 27
- “Configuring storage subsystems” on page 29
- “Creating direct and SAN-attached configurations” on page 35
- “Installing the RDAC failover driver” on page 36
- “Identifying devices on AIX hosts” on page 38
- “Verifying the installation and configuration” on page 40
- “Viewing and setting Object Data Manager (ODM) attributes” on page 48
- “Additional AIX configuration information” on page 55
- “Resolving disk array errors” on page 59
- “Redistributing logical drives in case of failure” on page 61

Hardware, software and host system requirements

This section lists the the following requirements:

- Hardware and firmware requirements for the DS4000 configuration
- Hardware and software requirements for the AIX host system

Note: For the latest installation requirements and user information about Storage Manager, AIX file sets, and hardware components, see the readme file that is located in the AIX DS4000 Installation CD. The most recent copy of the readme file, which supersedes this document, is maintained at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for AIX and click the link in the Current Version column.

DS4000 hardware and firmware requirements

Table 8 on page 24 lists the DS4000 Fibre Channel storage servers that you can use with DS4000 Storage Manager.

Table 8. DS4000 storage servers that are used with AIX systems

Product Name	Model	Minimum appware/snapware (firmware) and NVSRAM versions
IBM TotalStorage DS4100 SATA Storage Server	1724-100	For the most recent readme files, see the following Web site: www-1.ibm.com/servers/storage/support/disk/ Click the link for your storage server. When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for AIX and click the link in the Current Version column.
IBM TotalStorage DS4500 Fibre Channel Storage Server	1742-90U	
IBM TotalStorage DS4400 Fibre Channel Storage Server	1742-1RU	
IBM TotalStorage DS4300 Standard and Turbo Fibre Channel Storage Servers	1722-60U, 1722-60X	

Table 9 lists the DS4000 EXP storage expansion enclosures and other hardware that are supported with DS4000 Storage Manager.

Note: For important information about using the DS4000 EXP100 SATA storage expansion enclosure, see *IBM TotalStorage DS4000 EXP100 storage expansion enclosure Installation and User's Guide* (GC26-7601).

Table 9. Supported versions of hardware for AIX systems

Hardware	Product Name	Model	Minimum firmware version
storage expansion enclosures	IBM TotalStorage DS4000 EXP100 SATA storage expansion enclosure	1710-10U or 10X	9554
	IBM TotalStorage DS4000 EXP710Fibre Channel storage expansion enclosure	1740-710	9627
	IBM TotalStorage DS4000 EXP700 Fibre Channel storage expansion enclosure	1740-1RU or 1RX	9326
	IBM TotalStorage DS4000 EXP500 Fibre Channel storage expansion enclosure	3560-1RU or 1RX	9166

Table 9. Supported versions of hardware for AIX systems (continued)

Hardware	Product Name	Model	Minimum firmware version
Host Bus Adapters	IBM HA Emulex LP7000	FC 6227	3.30x1
	IBM HA Emulex LP9000	FC 6228	3.91a1
	IBM HA Emulex LP9802-F2	FC 6239	1.81x1
	IBM HA Emulex LP10000 (AIX 5.2H and 5.3 only)	FC 5716	1.90x4
Switches	Brocade switch	2109-S08 2109-S16	2.6.2a
	Brocade switch	2109-F16	3.1.2a
	Brocade switch	2109-M12	4.0.2c
	Brocade switch	2109-F32	4.2.0c
	Brocade switch	2109-F08	3.1.2a
	McData switch	2032-064, 2032-140	6.01.00 - 18
	McData switch	2031-016	6.01.00 - 18
	McData switch	2031-216, 2031-224	6.01.00 - 18
	InRange switch	2042-001	4.3.0
	Cisco switch	2062-D01, 2062-D07, 2062-T07	1.3.4a

AIX system requirements

This section lists the minimum hardware and software requirements for an AIX host in a DS4300, DS4400, or DS4500 storage subsystem that is managed by DS4000 Storage Manager.

Hardware requirements

The following hardware is required at a minimum:

Any RS/6000®, RS6000 SP or IBM @server pSeries™ or POWER™ server that supports FC 6227, FC 6228 or FC 6239 host bus adapters

Important: If you are using AIX 5.3 with a POWER4 system, contact your IBM service representative before installing AIX 5.3. Updated software is required.

Software requirements

The following AIX software is required:

- AIX 5.1, 5.2, or 5.3
- File sets associated with IBM RDAC driver installations that are listed in Table 11 on page 37, Table 12 on page 37, or Table 13 on page 37.

AIX restrictions

The following restrictions apply to AIX configurations:

SAN and connectivity restrictions

- AIX hosts (or LPARs) can support multiple host bus adapters (HBAs) and DS4000 devices. However, there is a restriction on the number of HBAs that can be connected to a single DS4000 storage server. You can configure up to two HBAs per partition and up to two partitions per DS4000 storage server. Additional HBAs can be added for additional DS4000 storage servers and other SAN devices, up to the limits of your specific server platform.
- Direct-attach configurations are restricted to single-initiator configurations only. You can only make one connection to each minihub on a DS4000 storage server.

Note: DS4300 Turbo storage servers do not have minihubs.

- Single HBA configurations are allowed, but each single HBA configuration requires that both controllers in the DS4000 be connected to the HBA, either through a switch or daisy-chain (direct-attach only). If they are connected through a switch, both controllers must be within the same SAN zone as the HBA.

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

- Single-switch configurations are allowed, but each HBA and DS4000 controller combination must be in a separate SAN zone.
- Other storage devices, such as tape devices or other disk storage, must be connected through separate HBAs and SAN zones.

Restrictions when booting up your system

- If you want to use a CD to install AIX on a DS4000 device on a partition with more than 32 LUNs, you must use a CD that installs AIX version 5.2H or later. You cannot use an AIX CD that installs an earlier version of AIX.

Therefore, if your AIX installation CD version is earlier than AIX 5.2H, and your system is booted from a DS4000 device, do not create more than 32 LUNs on the partition you boot from.

- When you boot your system from a DS4000 device, both paths to the DS4000 storage server must be up and running.
- The system cannot use path failover during the AIX boot process. Once the AIX host has started, failover operates normally.

Recommendation: It is recommended that you do not boot your system from any device that has one or more DS4000 EXP100 SATA storage expansion enclosures attached.

Partitioning restrictions

- The maximum number of partitions per AIX host, per DS4000 storage server, is two.
- All logical drives that are configured for AIX must be mapped to an AIX host group. For more information, see “Storage Partitioning: Defining an AIX host group” on page 33 and “Storage Partitioning premium feature” on page 21.
- On each controller, you must configure at least one LUN with an ID between 0 and 31 that is not a UTM or access logical drive.

Interoperability restrictions

- You cannot use F-RAID Manager.
- You cannot use dynamic volume expansion (DVE) on AIX 5.1.

Installing DS4000 Storage Manager

Use the following procedures, in the order listed, to install the DS4000 Storage Manager software on an AIX storage management station.

1. “Installing the client software”
2. “Installing SMagent (optional)” on page 28

Installing the client software

Install the client software in the following order:

1. DS4000 SMruntime software
2. DS4000 SMclient software

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

Prerequisites

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

- The AIX host on which you are installing SMruntime meets the minimum hardware and software requirements described in “Hardware, software and host system requirements” on page 23 and “Software requirements” on page 25.
- The correct file sets are present on the system.

Note: If the file sets are not present, follow the instructions in “Installing the RDAC failover driver” on page 36 to download them from the appropriate Web site. For lists of required file sets, see the following tables in “AIX system requirements” on page 25:

- **AIX 5.1:** Table 11 on page 37.
 - **AIX 5.2:** Table 12 on page 37.
 - **AIX 5.3:** Table 13 on page 37.
- The DS4000 Storage Manager files are available in a directory on the system.

Steps for installing SMruntime software

Modify the commands as needed for your installation. No restart is required during the installation process.

1. Install SMruntime by typing the following command:

```
# installp -a -d /path_name  
/SMruntime.aix-09.10.G5.00.bff SMruntime.aix.rte
```

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

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

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

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

Steps for installing SMclient software

Modify the commands as needed for your installation. No restart is required during the client installation process.

1. Install SMclient by typing the following command:

```
# installp -a -d /path_name/  
SMclient.aix-09.10.G5.01.bff SMclient.aix.rte
```

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

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

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

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

Installing SMagent (optional)

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

Prerequisites

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

- This is the AIX server that you identified as the storage management station.
- This machine meets the minimum hardware and software requirements described in “AIX system requirements” on page 25.
- SMruntime has been installed.

Note: SMagent is not supported with Storage Manager controller firmware version 5.42.xx.xx.

Steps for installing SMagent software

Modify these commands as needed for your installation. No restart is required during the client installation process.

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

1. Install SMagent by typing the following command:

```
# installp -a -d /path_name/  
SMagent.aix-09.10.G5.00.bff SMagent.aix.rte
```

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

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

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

Starting or stopping SMagent software

Use the following command to manually start SMagent:

```
# SMagent start
```


Use the following command to manually stop SMagent:

```
# SMagent stop
```

Configuring storage subsystems

If you have purchased a Storage Partitioning premium feature, ensure that the premium feature is enabled. For more information, see “Enabling your premium features” on page 21.

Complete the following procedures to configure DS4000 Storage Manager for an AIX system. This can be done from the DS4000 SMclient software running on either an AIX or a non-AIX system.

Adding storage subsystems to SMClient

Complete the following steps to add storage subsystems to DS4000 SMclient:

1. To set up the storage subsystem for AIX and DS4000 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 DS4000 storage controllers” on page 17 for information about assigning IP addresses to the controllers.

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

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

The Enterprise Management window opens.

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

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

Upgrading controller firmware and NVSRAM

The procedures in this section describe how to download and install DS4000 controller firmware and NVSRAM.

Before you begin: Read the information in the following sections before you begin the download procedures:

- “Checking current versions of firmware and NVSRAM”
- “Upgrading firmware and NVSRAM using concurrent download” on page 30
- “Staged download” on page 31

Checking current versions of firmware and NVSRAM

Versions of firmware and NVSRAM that are provided with this product might be later releases than those described in this document.

To check the firmware and NVSRAM versions that are currently installed on your system, complete the following steps:

1. In the Subsystem Management window, click **Storage Subsystem** → **View Profile**. The Storage Subsystem Profile window displays.
2. In the Storage Subsystem Profile window, click the Controllers tab to see the current versions of firmware and NVSRAM.
3. Click **Close** to close the Storage Subsystem Profile window.
4. To find whether the versions of firmware and NVSRAM that are currently on your storage subsystem are the latest versions, see the readme file at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for AIX and click the link in the Current Version column.

The latest versions are listed in the readme file.

5. If you find that you do not have the latest versions of the firmware and NVSRAM, click the **Back** button on your browser to return to the previous page. From there, you can download the correct versions to the host by clicking the appropriate link that is listed in the IBM DS4000 Controller Firmware matrix.

Upgrading firmware and NVSRAM using concurrent download

Depending on your system's current firmware and AIX device driver levels, you might be able to use *concurrent download*. Concurrent download is a method of downloading firmware to the controllers that does not require you to stop I/O to the controllers during the process.

Attention:

1. You cannot use concurrent firmware download if you change the default setting of the Object Data Manager (ODM) attribute *switch_retries*. The default is **5**. See "Viewing and setting Object Data Manager (ODM) attributes" on page 48 for more information.
2. If you do not have the correct firmware versions to use concurrent download, you must ensure that all I/O to the controllers is stopped before you upgrade the firmware or NVSRAM.

To find out whether you can use concurrent download, see Table 10.

Attention:

Table 10. Necessary criteria for concurrent download

Storage Manager controller firmware version:	Minimum AIX device driver file sets:
DS4300: 05.40.07.xx	devices.fcp.disk.array: 5.1.0.58
DS4400/DS4500: 05.40.06.xx	devices.fcp.disk.array: 5.2.0.17
All other DS4000 storage servers: 05.30.xx.xx, 05.40.xx.xx, 06.10.xx.xx	devices.fcp.disk.array: 5.3.0.0

Table 10. Necessary criteria for concurrent download (continued)

Storage Manager controller firmware version:	Minimum AIX device driver file sets:
<p>Attention: You cannot upgrade Storage Manager controller firmware version 5.4x.xx.xx to Storage Manager controller firmware version 6.1x.xx.xx using concurrent download. For example:</p> <ul style="list-style-type: none"> • You can only use concurrent download if the Storage Manager controller firmware version is already 6.1x.xx.xx, and you are downloading another version of 6.1x.xx.xx firmware. • You can only use concurrent download if the Storage Manager controller firmware version is already 5.4x.xx.xx, and you are downloading another version of 5.4x.xx.xx firmware. 	

Upgrading firmware using traditional or staged download

There are two methods for downloading the firmware to the controllers:

- “Traditional download”
- “Staged download”

Traditional download: The traditional download process takes significantly longer and must be done in one phase, rather than in two phases as with the staged controller firmware download. Therefore the preferred method of controller firmware download is the staged controller firmware download.

Staged download: With staged firmware download, the download process occurs in two distinct phases: firmware transfer and firmware activation. You can perform the time-consuming task of loading the firmware online so that it is functionally transparent to the application. You can then defer the activation of the loaded firmware to a convenient time.

Some reasons for activating firmware files at a later time include the following:

- **Time of day** - Activating can take a long time, so you might want to wait until I/O loads are lighter because the controllers will go offline briefly to activate the new firmware.
- **Type of package** - You might want to test the new firmware package on one storage subsystem before activating it on other storage subsystems.
- **Multiple controllers** - You can download controller firmware from the storage management software to all storage subsystem controllers for later activation.

Notes:

1. You cannot use staged download to upgrade Storage Manager controller firmware version 5.4x.xx.xx to version 6.1x.xx.xx. You can only use staged download if the Storage Manager controller firmware version is already 6.1x.xx.xx, and you are downloading another version of 6.1x.xx.xx firmware.
2. Staged controller firmware download is not supported on DS4400 storage servers.

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

Steps for upgrading the controller firmware: Complete the following steps to install a new firmware version on the controllers:

Before you begin:

1. Download the correct version of the firmware from the Web site that is listed in “Checking current versions of firmware and NVSRAM” on page 29. Place the file in a designated directory on the host system.
2. Ensure that you have management connections to both controllers, and that controllers are in an optimal state.
1. Check your system’s current firmware and device driver levels to determine whether you need to stop I/O to the controllers before you upgrade the firmware, using the information that is listed in Table 10 on page 30.
Attention: If you do not have the correct firmware and device driver levels for concurrent download, you *must* stop I/O to the controllers before you upgrade the firmware.
2. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.
3. In the Subsystem Management window, click **Advanced** —> **Maintenance** —> **Download** —> **Controller Firmware**. The Download Firmware window displays.
4. Click **Browse** to select the firmware file that you previously downloaded to a directory on the host.
5. If you want to perform a staged firmware download, select the check box next to “Transfer files but don’t activate them (activate later).” If you want to perform a traditional firmware download, do not select the check box, and the firmware will load immediately.
6. Click **OK**. The Confirm Download window displays.
7. Click **Yes** to start the download.
8. Close the Subsystem Management window and then reopen it to complete the firmware update.
9. Verify that the firmware was successfully installed.

Upgrading the NVSRAM

Complete the following steps to install a new version of NVSRAM on the controllers. NVSRAM will be activated immediately.

Before you begin:

1. Download the correct NVSRAM version from the Web site that is listed in “Checking current versions of firmware and NVSRAM” on page 29. Place the file in a designated directory on the host system.
2. Ensure that you have management connections to both controllers, and that controllers are in an optimal state.
1. Check your system’s current firmware and device driver levels to determine whether you need to stop I/O before upgrading the NVSRAM, using the information that is listed in Table 10 on page 30.
Attention: If you do not have the correct firmware and device driver levels for concurrent download, you *must* stop I/O to the controllers before you upgrade the NVSRAM.
2. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.
3. In the Subsystem Management window, click **Advanced** —> **Maintenance** —> **Download** —> **Controller NVSRAM**. The Download NVSRAM window displays.
4. Click **Browse** to select the NVSRAM file that you previously downloaded to a directory on the host.
5. Click **OK**. The Confirm Download window displays.

6. Click **Yes** to start the download.
7. Verify that the NVSRAM was successfully installed.

Storage Partitioning: Defining an AIX host group

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

For more information about storage partitioning, see “Storage Partitioning premium feature” on page 21.

This section contains the following procedures:

1. “Steps for defining an AIX host group”
2. “Steps for verifying that the host ports match the host” on page 34

Before you begin: Read the following information.

- Create the AIX host group at the storage subsystem level. Do not create host groups at the default group level.
Exception: If you are running a DS4300 configuration without partitioning enabled, you can use the default host group.
- In a cluster partition, perform logical drive mappings on the host group level so that all the hosts can see the same storage. In a normal partition, perform logical drive mappings on the host level.

Steps for defining 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**.

Note: Make sure that the storage subsystem is highlighted in the left panel of the Subsystem Management window. Do not click on Undefined Mappings.

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, then right-click and select **Define Host Port**.
 - c. Type or select the desired host port for the first HBA, and then change the host type to **AIX**. Click **Add**.

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

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

Steps for verifying that the host ports match the host

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

1. Type the following command:

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

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

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

2. Identify the fcs number of the HBA that is connected to the DS4000.
3. Type the following command:

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

where *fcs#* is the fcs number of the HBA that is connected to the DS4000. The network address number of the HBA is displayed, as in the following example:

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

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

4. Verify that the network address number matches the host port number that displays in host partition table of the DS4000 SMclient.
5. Repeat this procedure to verify the second host port.

Mapping LUNs to a storage partition

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

Before you begin: Read “DS4000 Storage Manager performance planning” on page 18 for information about planning the number of LUNs to map to each partition.

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

Note: It is recommended that you do not map at the host group level. If you do, every host can see all the LUNs.

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

Creating direct and SAN-attached configurations

DS4000 Storage Manager supports IBM DS4000 Fibre Channel Storage Servers in direct-attached AIX configurations or in a SAN environment through switches in AIX configurations.

Creating a direct-attached configuration

Before you begin: Ensure that:

- One or two AIX servers can be connected to the DS4000 storage servers.
- No external hubs are being used.
- Two-server DS4400 or DS4500 configurations require four host-side minihubs, each with exactly one Fibre Channel connection from each HBA to a minihub.

Note: DS4300 Turbo storage servers do not have minihubs.

- Single HBA configurations are allowed. If you have a single HBA in a direct-attached configuration, both controllers in the DS4000 must be connected to the HBA through a daisy-chain.

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

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

1. Connect the HBAs to each controller or minihub port of the DS4000 storage server.
2. Configure and verify the configuration, using the procedures that are described in “Identifying devices on AIX hosts” on page 38 and “Verifying the installation and configuration” on page 40.

Creating a SAN-attached configuration

Before you begin: Ensure that:

- Multiple Fibre Channel HBAs within the same server cannot “see” the same DS4000 controller port.
- The IBM Fibre Channel HBAs are isolated from each other if they are connected to the same switch that is connected to the same DS4000 controller port.
- Each Fibre Channel HBA and controller port must be in its own fabric zone, if they are connecting through a single Fibre Channel switch, such as 2109-F16.

Multiple DS4000 devices can be configured to the same set of Fibre Channel HBAs through a Fibre Channel switch.

Note: Single HBA configurations are allowed. If you have a single HBA in a SAN-attached configuration, both controllers in the DS4000 must be Fibre Channel connected to the HBA through a switch, and both controllers must be within the same SAN zone as the HBA. For more information about zoning and enabling zones, see Appendix E, “Connecting HBAs in a FC switch environment,” on page 165.

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

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

1. Connect the HBAs to the switch or switches.
2. Connect the DS4000 storage subsystems to the switch or switches.
3. Set the required zones on the Fibre Channel switch or switches, if applicable.

+ **Note:** For more information about zoning and enabling zones, see Appendix E,
+ “Connecting HBAs in a FC switch environment,” on page 165.

4. Configure and verify the configuration.
5. Use the **cfgmgr** command to enable the devices to log in to the switch.

Installing the RDAC failover driver

After you install the client software and configure your storage subsystems, use these instructions to install the appropriate DS4000 Storage Manager RDAC device driver.

You must install RDAC on all AIX hosts in your storage subsystem. This section describes how to check the current RDAC driver program driver version level, update the RDAC device driver, and verify that the RDAC update is complete.

Prerequisites

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

Installing the RDAC driver

The RDAC installation procedure that is described in “Steps for installing the RDAC driver” on page 37 refers to the tables listed in “AIX RDAC driver file sets.”

AIX RDAC driver file sets

The file sets that are listed in Table 11 on page 37, Table 12 on page 37, and Table 13 on page 37 must be the specified versions or later. Be sure to check the readme file for more recently-updated information.

A readme file that contains the latest installation and user information about the AIX file sets, storage management software, and hardware components is located on the AIX DS4000 Installation CD. You can find the most recent copy of the readme file, which supersedes this document, at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for AIX and click the link in the Current Version column.

Notes:

1. If you have 2102-F10 storage array devices connected to your system, do not install `devices.fcp.disk.array.rte` versions 5.1.0.58 or 5.2.0.17 or later. These versions do not support 2102-F10 devices. Contact your IBM technical support representative for assistance.
2. The AIX RDAC driver files are not included on the DS4000 installation CD. Either install them from the AIX Operating Systems CD, if the correct version is included, or download them from the following Web site:

<http://www-912.ibm.com/eserver/support/fixes/fcgui.jsp>

For more information, see “Installing the RDAC failover driver” on page 36.

Table 11. File sets required for AIX 5.1 RDAC

PTF file sets	Version
devices.fcp.disk.array.diag	5.1.0.51
devices.fcp.disk.array.rte	5.1.0.60
devices.common.IBM.fc.rte	5.1.0.50
devices.pci.df1000f7.com	5.1.0.58
devices.pci.df1000f7.rte	5.1.0.15
devices.pci.df1000f9.rte	5.1.0.35

Table 12. File sets required for AIX 5.2 RDAC

PTF file sets	Version
devices.fcp.disk.array.diag	5.2.0.11
devices.fcp.disk.array.rte	5.2.0.31
devices.common.IBM.fc.rte	5.2.0.10
devices.pci.df1000f7.com	5.2.0.31
devices.pci.df1000f7.rte	5.2.0.10
devices.pci.df1000f9.rte	5.2.0.10

Table 13. File sets required for AIX 5.3 RDAC

PTF file sets	Version
devices.fcp.disk.array.diag	5.3.0.0
devices.fcp.disk.array.rte	5.3.0.0
devices.common.IBM.fc.rte	5.3.0.0
devices.pci.df1000f7.com	5.3.0.0
devices.pci.df1000f7.rte	5.3.0.0
devices.pci.df1000f9.rte	5.3.0.0

Steps for installing the RDAC driver

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

Before you begin: Check the file set versions that are listed in “AIX system requirements” on page 25 or in the readme file, to verify whether your AIX systems have the most recent RDAC file sets. You only need to perform this installation on AIX systems that do not have the most recent RDAC file sets.

1. Download the most recent file sets from the following Web site:

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

2. Verify that the correct version of the software was successfully installed by typing the following command:

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

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

- For each AIX host that is to be attached to a DS4000 storage subsystem, repeat the above steps in order to install and verify the RDAC driver.

After you have finished installing RDAC, you can perform the initial device discovery, which is described in the next section.

Identifying devices on AIX hosts

The RDAC driver creates the following devices that represent the DS4000 storage subsystem configuration, as shown in Figure 8:

- dar** The disk array router (dar) device represents the entire array, including the current and the deferred paths to all LUNs (hdisks).
- dac** The disk array controller (dac) devices represent a controller within the storage subsystem. There are two dacs in the storage subsystem.
- hdisk** Each hdisk device represents an individual LUN on the array.
- utm** The universal transport mechanism (utm) device is used only with in-band management configurations, as a communication channel between the SMagent and the DS4000.

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

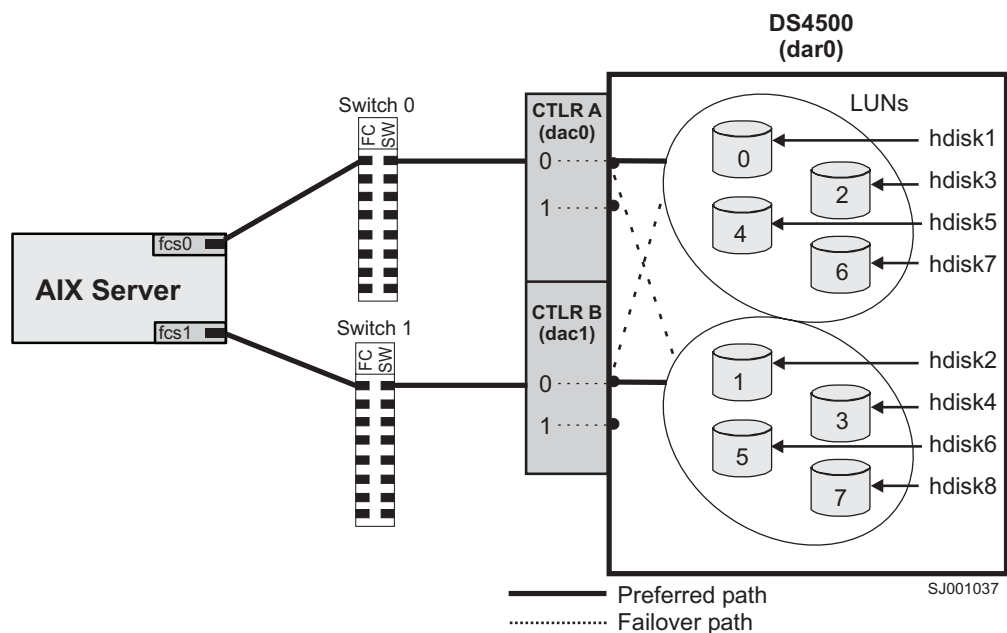


Figure 8. DS4000/AIX system configuration

Performing initial device discovery

Complete these steps to perform the initial device discovery.

Before you begin: Ensure that the DS4000 storage subsystem has been set up, LUNs have been assigned to the host, and the RDAC driver has been installed.

1. Type the following command to probe for the new devices:

```
# cfgmgr -v
```

Note: In a SAN configuration, the devices do not log into the SAN switch until you run **cfgmgr**.

2. Type the following command:

```
# lsdev -Cc disk
```

3. Examine the output of the **lsdev -Cc disk** command to ensure that the RDAC software recognizes the DS4000 logical drives, as shown in the following list:
 - Each DS4300 logical drive is recognized as a “1722 (600) Disk Array Device”.
 - Each DS4400 logical drive is recognized as a “1742 (700) Disk Array Device”.
 - Each DS4500 logical drive is recognized as a “1742-900 Disk Array Device”.

The following example shows the output of the **lsdev -Cc disk** command for a set of DS4500 LUNs:

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

Important: You might discover that the configuration process has created two dacs and two dars on one DS4000 subsystem. This situation can occur when your host is using a partition that does not have any associated LUNs. When that happens, the system cannot associate the two dacs under the correct dar. If there are no LUNs, the system generates two dacs as expected, but it also generates two dars.

The following list shows the most common causes:

- You create a partition and attach the LUNs to it, but you do not add the host ports to the partition. Therefore, the host ports remain in the default partition.
- You replace one or more HBAs, but do not update the worldwide name (WWN) of the partition for the HBA.
- You switch the DS4000 from one set of HBAs to another as part of a reconfiguration, and do not update the WWNs.

In each of these cases, resolve the problem, and run **cfgmgr** again. The system removes the extra dar, or moves it from the Available state to the Defined state. (If the system moves the dar into the Defined state, you can then delete it.)

Note: When you perform the initial device identification, the Object Data Manager (ODM) attributes of each device are updated with default values. In most cases and for most configurations, the default values are satisfactory.

However, there are some values that can be modified for maximum performance and availability. See “Viewing and setting Object Data Manager (ODM) attributes” on page 48.

See “Using the `lsattr` command to view ODM attributes” on page 54 for information about using the `lsattr` command to view attribute settings on an AIX system.

Verifying the installation and configuration

After you have performed the initial device identification, complete the following procedures to verify that all of your DS4000 device names and paths are correct and that AIX recognizes your dars, dacs, and hdisks.

- “Identifying the controller ID numbers”
- “Identifying device names and bus numbers” on page 41
- “Identifying logical drives by operating system device names” on page 46
- “Identifying FlashCopy logical drive types” on page 47

Tip: Verify your installation and configuration before you mount your file systems and install your applications.

Identifying the controller ID numbers

Choose one of the following options to identify the controller ID number of a dac:

- “Using the `lsattr -El` command”
- “Using the SMClient Controller Properties window”

Using the `lsattr -El` command

1. Type the `lsattr -El` command, as follows:

```
# lsattr -El dacname
```

where *dacname* is the name of the dac whose controller ID number you want to identify.

2. Find the value of the `controller_SN` attribute, as shown in the following example. This value is the controller ID number of dac1.

```
# lsattr -El dac1
GLM_type          low          GLM type          False
alt_held_reset    no           Alternate held in reset  False
cache_size        1024        Cache Size in MBytes  False
controller_SN    1T14148774 Controller serial number False
ctrl_type         1742-0900   Controller Type      False
location          location     Location Label       True
lun_id            0x0         Logical Unit Number  False
node_name         0x200400a0b80f434d FC Node Name        False
passive_control    no           Passive controller   False
scsi_id           0x11a00     SCSI ID              False
utm_lun_id        none        Logical Unit Number  False
ww_name           0x200500a0b80f434e World Wide Name      False
```

Using the SMClient Controller Properties window

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

2. In the Controller Properties window, find the value of the Serial number attribute. In the example shown in Figure 9, the controller in slot A has the controller ID number (serial number) 1T14148774, and is represented by dac1.

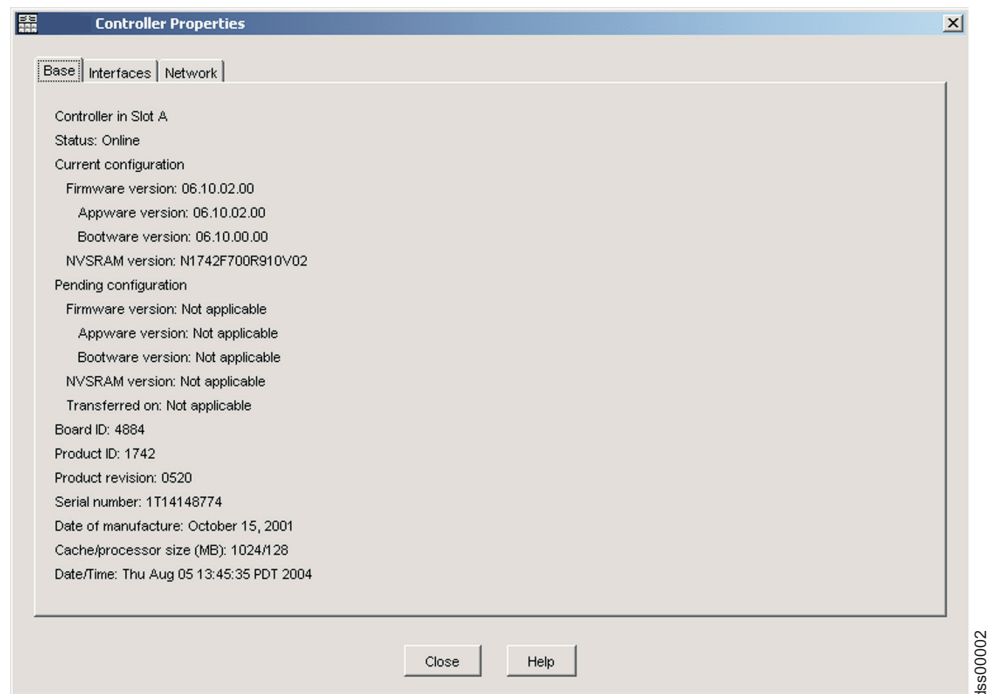


Figure 9. Controller Properties window

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 DS4000 device, and then from each path to its associated logical drive. As mentioned previously in this chapter, the RDAC software uses `dars`, `dacs`, and `hdisks` to represent the DS4000 storage subsystem. `dars` represent the disk array routers; `dacs` represent the disk array controllers, and `hdisks` represent the logical drives, or DS4000 volumes.

AIX provides the following commands to help you determine the DS4000 configuration, and to get information about device names and bus numbers:

lsdev

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

lsattr

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

fget_config

Displays controllers and `hdisks` that are associated with a specified DS4000 (`dar`).

The **fget_config** command shows the current state and logical drive (`hdisk`) ownership.

There are several ways to correlate a system's configuration and monitor the state of DS4000 storage subsystems.

Using the `lsdev` command

The examples in this section show you how to use the `lsdev` command to show the status of `dars`, `dacs` and `hdisks`.

lsdev example 1: This example shows the use of the `lsdev` command to view the status of the `dar`, which represents a DS4000 storage subsystem.

In the example, `dar0` is a machine type 1742, which is a DS4400. It is in the Available state, which is the state at the time when the device was last configured by AIX.

```
# lsdev -C |grep dar0
dar0      Available          1742      (700) Disk Array Router
```

lsdev example 2: This example shows the use of the `lsdev` command to view the status of two `dacs`, which represent the DS4000 storage subsystem controllers.

In the example, as in the previous example, a DS4400 is shown in an Available state. Each `dac` has its own location code, or *path*, which are represented by the values 11-08-01 and 14-08-01. Each AIX system has its own set of location codes that describe the internal path of that device, including bus and host-adaptor locations.

See the service manual for your system type to identify device locations.

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

lsdev example 3: This example shows the use of the `lsdev` command to view the status and location codes of eight DS4400 `hdisks`.

In the example, note that the location codes of the odd-numbered `hdisks` correspond to the location code of `dac0` that is listed in the previous example, and that the location codes of the even-numbered `hdisks` correspond to the location code of `dac1`. This means that the preferred I/O path for `hdisks` 1, 3, 5, and 7 is through `dac0`, and the failover path is through `dac1`. Conversely, the preferred path for `hdisks` 2, 4, 6, and 8 is through `dac1`, and failover path is through `dac0`.

```
# lsdev -Cc |grep hdisk
hdisk0    Available 40-60-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1    Available 11-08-01      1742      (700) Disk Array Device
hdisk2    Available 14-08-01      1742      (700) Disk Array Device
hdisk3    Available 11-08-01      1742      (700) Disk Array Device
hdisk4    Available 14-08-01      1742      (700) Disk Array Device
hdisk5    Available 11-08-01      1742      (700) Disk Array Device
hdisk6    Available 14-08-01      1742      (700) Disk Array Device
hdisk7    Available 11-08-01      1742      (700) Disk Array Device
hdisk8    Available 14-08-01      1742      (700) Disk Array Device
```

lsdev example 4: This example shows the use of the `lsdev` command with the `grep` command to search for the first location code two numbers. In this way, you

can determine the full preferred path from hdisk to Fibre Channel fabric.

```
# 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    1742      (700) Disk Array Controller
hdisk1    Available 11-08-01    1742      (700) Disk Array Device
hdisk3    Available 11-08-01    1742      (700) Disk Array Device
hdisk5    Available 11-08-01    1742      (700) Disk Array Device
hdisk7    Available 11-08-01    1742      (700) Disk Array Device
hdisk8    Available 11-08-01    1742      (700) Disk Array Device
```

Using the `fget_config` command

You can use the `fget_config` command to display current configuration information for a DS4000 subsystem.

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

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

- **A** Displays output for all the configured dars in the subsystem. If you use this parameter, do not specify a dar name.
- **l** Displays output only for the dar that you specify.
- **v** Displays more information about the dar or dars, such as the user array names that were created when you configured the DS4000 subsystem.

fget_config example 1: This example shows a sample output from the `fget_config -l dar3` command. A dac can be Active, Passive, or Offline. Only dacs that are in Active state are used for I/O. In the output, you can see that dar3 consists of dac6 and dac9.

```
# fget_config -l dar3
dac6 ACTIVE dac9 ACTIVE
hdisk26 dac9
hdisk27 dac6
hdisk28 dac9
hdisk29 dac6
hdisk30 dac9
hdisk31 dac6
hdisk32 dac9
hdisk33 dac9
hdisk83 dac6
```

fget_config example 2: This example shows a sample output from the `fget_config -vl dar3` command. Note that using the `-v` parameter causes the user array name and the LUN names to be displayed:

```
# fget_config -v1 dar3

User array name = 'D59 DS4400 #1'
dac6 ACTIVE dac9 ACTIVE

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

fget_config example 3: This example shows a sample output from the **fget_config -A** command. Note that using the **-A** parameter causes all the dars to be displayed:

```
# fget_config -A

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

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

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

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

fget_config example 4: This example shows a sample output from the **fget_config -Av** command. Note that using the **-v** parameter causes the user array names and the LUN names to display:


```

# fget_config -Av

---dar0---
User array name = 'D59S Flute-2
(DS4500) #1'
dac3 ACTIVE dac0 PASSIVE

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

---dar1---
User array name = 'D59 DS4500 #1'
dac1 ACTIVE dac4 ACTIVE

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

---dar2---
User array name = 'D59S DS4500900 #2'
dac2 ACTIVE dac5 ACTIVE

Disk    DAC    LUN Logical Drive
utm     DAC    LUN Logical Drive
utm     31
hdisk25 dac5    0  1

---dar3---
User array name = 'D59 DS4500 #1'
dac6 ACTIVE dac9 ACTIVE

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

```

Using the lsattr command

You can also use the **lsattr** command to find out which controller owns a particular hdisk.

The following example shows portions of the outputs from several **lsattr** commands.

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.

```

# lsattr -El dac0
GLM_type      low                GLM type                False
alt_held_reset no                Alternate held in reset False
cache_size    1024                Cache Size in MBytes    False
controller_SN 1T23559428          Controller serial number False
ctrl_type     1742-0900           Controller Type          False
location      Location Label        True
lun_id        0x0                Logical Unit Number      False
node_name     0x200400a0b80f434d FC Node Name            False
passive_control no                Passive controller       False
scsi_id      0x11100           SCSI ID                  False
utm_lun_id    none                Logical Unit Number      False
ww_name       0x200400a0b80f434e World Wide Name         False

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

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

```

Identifying logical drives by operating system device names

The **lsattr** command provides detailed information about a logical drive, including information that allows you to map the system device name to the logical drive on the DS4000 storage subsystem.

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

```

# lsattr -El hdisk4
pvid          none                Physical volume identifier False
q_type        simple            Queuing Type            False
queue_depth   30                Queue Depth             True
reserve_lock  yes                RESERVE device on open  True
write_cache   yes                Write Caching enabled   True
size         6062            Size in Mbytes          False
raid_level   0                RAID Level              False
rw_timeout    30                Read/Write Timeout value True
reassign_to   120               Reassign Timeout value  True
scsi_id       0x11000           SCSI ID                  False
lun_id       0x0002000000000000 Logical Unit Number      False
cache_method  fast_write        Write Caching method    True
prefetch_mult 0                Multiple of blocks to prefetch on read True
ieee_volname 600A0B80000F14AF000003D140C046A2 IEEE Unique volume name False

```

In the example, you can see that hdisk4 is a 6 GB LUN of type RAID 0, with a LUN ID of 2, and an IEEE volume name of 600A0B8000F14AF000003D140C046A2.

You can make a more exact correlation using the distinctive *ieee_volname* attribute. The value of this attribute on the AIX host is the same as the value of the *Logical Drive ID* on the DS4000 subsystem. You can see the *Logical Drive ID* in the Change Mapping window of DS4000 Storage Manager, which is shown in Figure 10.

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

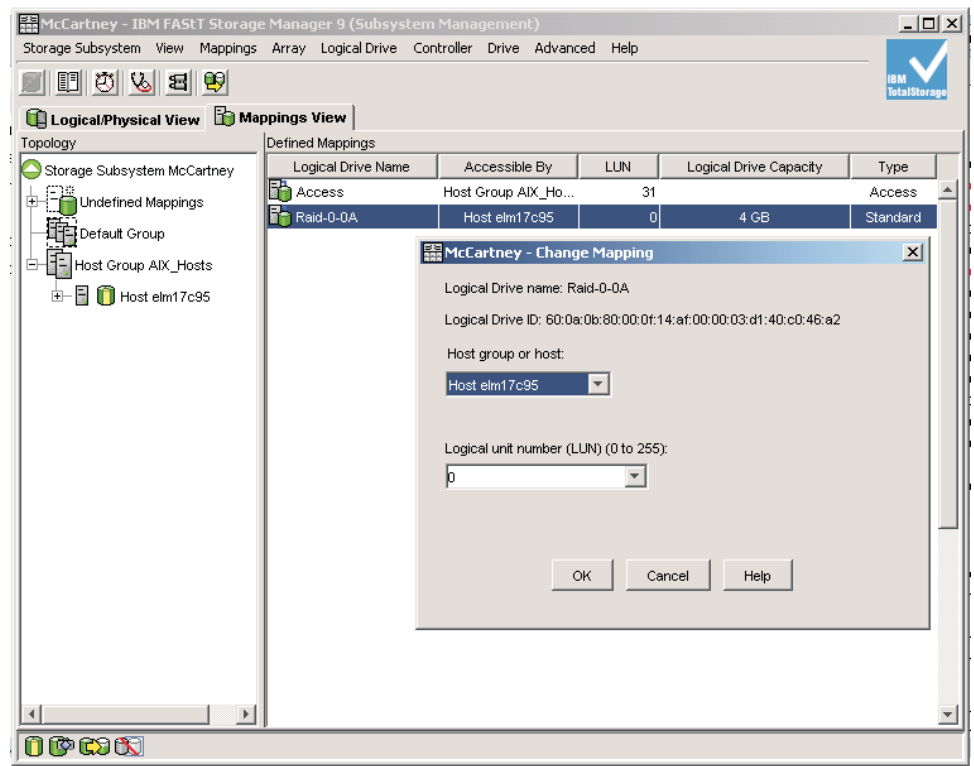


Figure 10. Change Mappings window

You can also identify the system disk name by viewing the *Subsystem ID (SSID)* attribute in the Logical Drive Properties window of DS4000 Storage Manager. The value of the *Subsystem ID (SSID)* attribute on the DS4000 subsystem is the same as the value of the *lun_id* attribute on the AIX host.

To open the Logical Drive Properties window, right-click the name of the logical drive in the Logical/Physical view of the Subsystem Management window and select **Properties**.

Identifying FlashCopy logical drive types

Complete the following steps to identify FlashCopy logical drive types:

1. Run the following command to list available logical drives:

```
# lsdev -Cc disk
```

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

```
[root@hostname] / # lsdev -Cc disk |pg
hdisk0 Available 40-60-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 40-60-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 11-08-01 1742-900 (900) Disk Array Device
hdisk3 Available 11-08-01 1742-900 (900) Disk Array Device
hdisk4 Available 11-08-01 1742-900 (900) Disk Array Snapshot
```

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

2. Run the following command to display the state of each controller in a DS4000 array:

```
# fget_config -Av
```

A list displays, as shown in the following example:

```
[root@hostname] / # fget_config -Av
---dar0---
User array name = 'DS4000 System 1'
dac0 ACTIVE dac1 ACTIVE
dac0-hdisk2 Raid-0-0B
dac1-hdisk3 Raid-0-1B
dac1-hdisk4 Raid-3-0A-1 Snapshot Volume
```

The FlashCopy logical drive is identified by the term Snapshot Volume.

For more information about the **lsdev** and **fget_config** commands, see “Identifying device names and bus numbers” on page 41.

Viewing and setting Object Data Manager (ODM) attributes

Some of the ODM attributes are for information purposes only. These information-only attributes show how the DS4000 storage server is configured, or its current state. You can modify other attributes using SMIT or by using the AIX **chdev -p** command.

Attribute definitions

The following tables list definitions and values of the ODM attributes for dars, dacs and hdisks:

- Table 14 on page 49: *Attributes for dar devices*
- Table 15 on page 50: *Attributes for dac devices*
- Table 16 on page 51: *Attributes for hdisk devices*

Notes:

1. Attributes with True in the Changeable column can be modified from their default settings.
2. Attributes with False in the Changeable column are for informational or state purposes only. However, some attributes with False in the Changeable column can be modified using DS4000 Storage Manager.

3. The **lsattr -el** command is another way to determine which attributes can be modified. Attributes that can be modified display True in the last column of the **lsattr -el** output. You can also display the default values by using the **lsattr -dl** command. See “Using the lsattr command to view ODM attributes” on page 54 for more information.

Table 14. Attributes for dar devices

Attribute	Definition	Changeable (T/F)	Possible value
<i>act_controller</i>	List of controllers in the active state at the time of configuration.	False	Set at configuration time by the RDAC software.
<i>all_controller</i>	List of controllers that comprise this array; usually there are two dac devices.	False	Set at configuration time by the RDAC software.
<i>held_in_reset</i>	Name of the controller that was in the held-in-reset state at the time of configuration, or none if no controllers were in that state.	True	Set at configuration time by the RDAC software. Should not be changed.
<i>load_balancing</i>	Indicator that shows whether load balancing is enabled (yes) or disabled (no); see the definition of the <i>balance_freq</i> attribute for more information.	True	Yes or No. Attention: You should only set the <i>load_balancing</i> attribute to yes in single-host configurations.
<i>autorecovery</i>	Indicator that shows whether the device returns the array to dual-active mode when it detects proper operation of both paths and controllers (yes) or not (no).	True	Yes or No. See restrictions on use.
<i>hlthchk_freq</i>	Number that specifies how often health checks are performed, in seconds.	True	1 - 9999. Should not be changed
<i>aen_freq</i>	Number that specifies how often polled AEN checks are performed, in seconds.	True	1 - 9999. Should not be changed
<i>balance_freq</i>	If <i>load_balancing</i> is enabled, number that specifies how often the system performs load-balancing on the array, in seconds.	True	1 - 9999 - should not be changed

Table 14. Attributes for dar devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<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 DS4000 configuration.
<i>cache_size</i>	Cache size for both controllers, in megabytes; 0 if the sizes do not match.	False	512 or 1024. Set by DS4000.
<i>switch_retries</i>	Number that specifies how many times to retry failed switches, in integers.	True	0 - 255. Default: 5 For most configurations, the default is the best setting. If you are using HACMP, it can be helpful to set the value to 0 . Attention: You cannot use concurrent firmware download if you change the default setting.

Table 15. 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 DS4000 configuration.
<i>alt_held_reset</i>	Indicator that shows whether the alternate controller was in the held-in-reset state at the time of configuration (yes) or not (no).	False	Yes or No. State of DS4000 configuration.
<i>controller_SN</i>	Serial number of this controller.	False	Set by DS4000.
<i>ctrl_type</i>	Type of array this controller belongs to. A value of value of 1722 indicates DS4300; a value of 1742 indicates DS4400; a value of 1742-900 indicates DS4500.	False	1742, 1722, 1742-900. Set by DS4000.

Table 15. Attributes for dac devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<i>cache_size</i>	Cache size of this controller, in megabytes.	False	512, 1024. Set by DS4000.
<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 DS4000.
<i>utm_lun_id</i>	Logical unit number of this controller, or none if UTM (access logical drives) is not enabled.	False	0 - 31. Set by DS4000 Storage Manager.
<i>node_name</i>	Name of the Fibre Channel node.	False	Set by DS4000
<i>location</i>	User-defined location label for this controller; the system does not use this value.	True	Set by DS4000 Storage Manager.
<i>ww_name</i>	Fibre Channel worldwide name of this controller.	False	Set by DS4000.
<i>GLM_type</i>	GLM type used for this controller.	False	High or Low. Set by DS4000.

Table 16. Attributes for hdisk devices

Attribute	Definition	Changeable (T/F)	Possible value
<i>pvid</i>	AIX physical volume identifier, or none if not set.	False	Set by AIX.
<i>q_type</i>	Queueing type for this device; must be set to simple .	False	Set by AIX. Must be "simple".
<i>queue_depth</i>	Number that specifies the depth of the queue based on system configuration; reduce this number if the array is returning a BUSY status on a consistent basis.	True	1 - 64 Note: See "Setting the queue depth for hdisk devices" on page 53 for important information about setting this attribute.
<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.

Table 16. Attributes for hdisk devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<i>write_cache</i>	Indicator that shows whether write-caching is enabled on this device (yes) or not (no); see the definition of the <i>cache_method</i> attribute for more information.	False	Yes or No.
<i>size</i>	Size of this logical drive.	False	Set by DS4000.
<i>raid_level</i>	Number that specifies the RAID level of this device.	False	0, 1, 3, 5. Set by DS4000 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 - 255. Set by DS4000 Storage Manager.

Table 16. Attributes for hdisk devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<i>cache_method</i>	<p>If <i>write_cache</i> is enabled, the write-caching method of this array; set to one of the following:</p> <ul style="list-style-type: none"> • default. Default mode; the word "default" is not seen if <i>write_cache</i> is set to yes. • fast_write. Fast-write (battery-backed, mirrored write-cache) mode. • fw_unavail. Fast-write mode was specified but could not be enabled; write-caching is not in use. • fast_load. Fast-load (non-battery-backed, non-mirrored write-cache) mode. • fl_unavail. Fast-load mode was specified but could not be enabled. 	False	Default, fast_write, fast_load, fw_unavail, fl_unavail.
<i>prefetch_mult</i>	Number of blocks to be prefetched into read cache for each block read.	False	0 - 100.
<i>ieee_volname</i>	IEEE unique logical drive name identifier for this logical drive.	False	Set by DS4000.

Setting the queue depth for hdisk devices

Setting the *queue_depth* attribute to the appropriate value is important for system performance. If you have a large DS4000 configuration with many logical drives and hosts attached, use this setting for high availability.

This section shows how to set the *queue_depth* attribute using the **chdev -l** command. To make the attribute changes permanent in the Customized Devices object class, use **chdev -P**.

Attention: If you set this attribute incorrectly, it could result in the loss of file systems.

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

• **Storage Manager controller firmware versions 06.1x.xx.xx or 05.4x.xx.xx**

On DS4000 storage systems that are running Storage Manager controller firmware versions 05.4x.xx.xx, 06.1x.xx.xx, or later, use the following formula to determine the maximum queue depth:

$$2048 / (\text{number-of-hosts} * \text{LUNs-per-host})$$

For example, a system with four hosts, each with 32 LUNs, would have a maximum queue depth of 16:

$$2048 / (4 * 32) = 16$$

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

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

• **Storage Manager controller firmware version 05.30.xx.xx**

On DS4000 storage systems that are running Storage Manager controller firmware version 05.30.xx.xx or earlier, use the following formula to determine the maximum queue depth:

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

For example, a system with four hosts, each with 32 LUNs, would have a maximum queue depth of 4:

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

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

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

Using the **lsattr** command to view ODM attributes

To view the Object Data Manager (ODM) attribute settings for dars, dacs, and hdisks, use the **lsattr** command, as follows:

- To view the default settings, type **lsattr -DI**.
- To view the attributes that are currently set on the system, type **lsattr -EI**.

The following **lsattr -EI** output examples show ODM attribute settings for a dar, a dac and an hdisk.

Example 1: Attribute settings for dar0:

```
# lsattr -EI dar0
act_controller dac0,dac1 Active Controllers          False
aen_freq      600      Polled AEN frequency in seconds          True
all_controller dac0,dac1 Available Controllers          False
autorecovery  no         Autorecover after failure is corrected      True
balance_freq  600      Dynamic Load Balancing frequency in seconds True
cache_size    128      Cache size for both controllers             False
fast_write_ok yes       Fast Write available                        False
held_in_reset none     Held-in-reset controller                    True
hlthchk_freq  600      Health check frequency in seconds          True
load_balancing no       Dynamic Load Balancing                      True
switch_retries 5       Number of times to retry failed switches    True
```

Example 2: Attribute settings for dac0:

```
# lsattr -El dac0
GLM_type          low          GLM type          False
alt_held_reset    no           Alternate held in reset False
cache_size        128         Cache Size in MBytes False
controller_SN     1T24594458 Controller serial number False
ctrl_type         1722-600   Controller Type   False
location          location    Location Label    True
lun_id            0x0         Logical Unit Number False
node_name         0x200200a0b80f14af FC Node Name      False
passive_control   no          Passive controller False
scsi_id           0x11000    SCSI ID           False
utm_lun_id        0x001f000000000000 Logical Unit Number False
ww_name           0x200200a0b80f14b0 World Wide Name   False
```

Example 3: Attribute settings for hdisk174:

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

Additional AIX configuration information

This section contains the following procedures:

- “Disabling DS4000 cache mirroring”
- “Using fast I/O failure for Fibre Channel devices” on page 56
- “Using dynamic tracking of Fibre Channel devices” on page 57
- “Using dynamic capacity expansion and dynamic volume expansion” on page 57

Disabling DS4000 cache mirroring

Cache mirroring is enabled on the DS4000 storage server by default. However, you might want to disable this function in certain environments where high performance is a critical factor (such as streaming media applications).

For most applications, cache mirroring should remain enabled.

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

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

- On AIX 5.1, 5.2, and 5.3, cache mirroring is *not* automatically re-enabled the next time the system is restarted or when you run the **cfgmgr** command.

Steps for disabling cache mirroring

In DS4000 Storage Manager, complete the following steps to disable cache mirroring.

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

1. In the Subsystem Management window, right-click the logical drive on which you want to disable cache mirroring, and select **Change** → **Cache Settings**.
2. In the Change Cache Settings window, clear the *Enable write caching with mirroring* check box.
3. Click **OK**.

Notes:

1. When a LUN is opened that is running with write cache enabled and cache mirroring disabled, an FCP array warning message displays. The warning displays again every 24 hours until cache mirroring is enabled again. See FCP_ARRAY_ERR25, which is described in “Resolving disk array errors” on page 59.

Using fast I/O failure for Fibre Channel devices

I/O failover occurs when a Fibre Channel adapter driver detects a link failure on the fabric between the switch and the DS4000 storage server. You can change the failover characteristics by setting the fscsi device attribute **fc_err_recov** to one of the following settings.

fast_fail Enables fast I/O failure.

If the Fibre Channel adapter driver detects that a link was lost between the switch and the DS4000, it waits a few seconds to allow the fabric to stabilize. If **fast_fail** is specified, the adapter then begins failing all I/O at the adapter driver if it detects that the device is not on the fabric. Any new I/O or future retries of the failed I/O are failed immediately.

Fast I/O failure can be useful in multipath configurations. It can decrease the I/O fail times due to the loss of a link between the storage device and the switch, and can allow faster failover to alternate paths.

delayed_fail Default setting.

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

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

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

```
chdev -l fscsi0 -a fc_err_recov=fast_fail
```

Notes:

1. The **fast_fail** attribute only affects failover that occurs between the switch and the DS4000 storage server. It does not affect failover that occurs between the host and the switch.
2. Set the **fast_fail** attribute on each HBA that is configured to the DS4000 storage server.
3. You can use fast I/O failure only in a SAN environment. You cannot use it in a direct-attach environment.

Using dynamic tracking of Fibre Channel devices

You can dynamically track Fibre Channel devices, which allows the dynamic movement of a Fibre Channel path between the fabric switch and the DS4000 subsystem by suspending I/O for 15 seconds while the move occurs.

Exception: You can only use dynamic tracking on AIX 5.2 and 5.3.

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

yes Enables dynamic tracking.

If dynamic tracking is enabled, the Fibre Channel adapter detects when the Fibre Channel node port ID of a device changes. It reroutes the traffic that is destined for that device to the new worldwide port name (WWPN) while the devices are still online.

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

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

no Default setting.

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

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

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

Notes:

1. Set the **dyntrk** attribute on each HBA that is configured to the DS4000 storage server.
2. You can use dynamic tracking only in a SAN environment. You cannot use it in a direct-attach environment.

Using dynamic capacity expansion and dynamic volume expansion

Dynamic volume expansion (DVE) is dynamic on the DS4000, but it requires manual intervention for AIX to recognize the new logical drive capacity. This section explains how to use DVE with AIX.

Exception: You can only use DVE on AIX 5.2 and 5.3.

Before you begin: Ensure that there is available free capacity within the array. You can check free capacity availability using DS4000 Storage Manager, in the Logical/Physical view of the Subsystem Management window. If there is not enough free capacity, and extra drives are available, perform a dynamic capacity expansion (DCE) operation before you perform the DVE operation. A DCE operation increases the capacity of the array by adding physical disks.

Performing a dynamic capacity expansion operation

Complete the following steps to perform a DCE:

1. From the Logical/Physical view of the Subsystem Management window, select **Array —> Add Free Capacity (Drives)**.
2. From the Logical/Physical window again, select **Logical Drive —> Increase Capacity**. The Increase Logical Drive Capacity — Additional Instructions window opens.
3. Read the additional instructions and click **OK**. The Increase Logical Drive Capacity window opens.
4. Type the amount that you want to increase the logical drive, and click **OK**.

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

Tip: If the DS4000 is busy, the process might take several hours to complete.

After you have ensured that there is sufficient free capacity within the array, you can perform DVE, as described in the next topic.

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

Performing a dynamic volume expansion operation

Complete the following steps to perform a DVE on the AIX host. For more information about expanding Logical Volume Manager (LVM) logical drives, see the man page for the **chvg** command .

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

```
umount mount_point
```

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

2. Change the logical drive group by typing the following command:

```
chvg -g logical drive group name
```

where the *logical drive group name* is the name of the associated LVM logical drive group.

Note: If you see a warning that the logical drive group cannot be imported to AIX 5.1 or lower, ignore the warning. It is an error. You can import the logical drive group to AIX 5.1 (or lower.)

3. Mount the file systems.

Notes:

1. You might have to run the **varyoffvg** command, and then the **varyonvg** command, on the logical drive group for LVM to see the size change on the disk. However, the new capacity is available for use by the operating system.

2. You cannot resize while the logical drive group is activated in classic or enhanced concurrent mode.
3. You cannot resize the root logical drive group.

Resolving disk array errors

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

You can also check your DS4000 Storage Manager Major Event log (MEL) to find out whether there is any correlation between the host, SAN, and DS4000 storage subsystem.

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

Note: For more information about troubleshooting, see the *IBM TotalStorage DS4000 Fibre Channel Hardware Maintenance Manual and Problem Determination Guide*.

- **FCP_ARRAY_ERR1** *ARRAY OPERATION ERROR*
A permanent hardware error involving the disk array media.
- **FCP_ARRAY_ERR2** *ARRAY OPERATION ERROR*
A permanent hardware error.
- **FCP_ARRAY_ERR3** *ARRAY OPERATION ERROR*
A permanent error detected by the array adapter.
- **FCP_ARRAY_ERR4** *ARRAY OPERATION ERROR*
A temporary error within the array, communications, adapter, and so on.
- **FCP_ARRAY_ERR5** *UNDETERMINED ERROR*
An undetermined error has occurred.
- **FCP_ARRAY_ERR6** *SUBSYSTEM COMPONENT FAILURE*
A degradation condition has occurred other than a disk drive.
- **FCP_ARRAY_ERR7** *CONTROLLER HEALTH CHECK FAILURE*
A health check on the passive controller has failed.
- **FCP_ARRAY_ERR8** *ARRAY CONTROLLER SWITCH*
One array controller has become unavailable, so I/O has moved to the other controller.
- **FCP_ARRAY_ERR9** *ARRAY CONTROLLER SWITCH FAILURE*
An array controller switch has failed.
- **FCP_ARRAY_ERR10** *ARRAY CONFIGURATION CHANGED*
A logical unit has been moved from one controller to the other (most likely by the action of an alternate host).
- **FCP_ARRAY_ERR11** *IMPROPER DRIVE TYPE FOR DUAL ACTIVE MODE*
This error should not be possible on the 2102 array, and exists for history reasons only. FCP_ARRAY_ERR11 might be reused for a different error in the future.
- **FCP_ARRAY_ERR12** *POLLED AEN FAILURE*
An automatic error notification has failed.
- **FCP_ARRAY_ERR13** *ARRAY INTER-CONTROLLER COMMUNICATION FAILURE*

The controllers are unable to communicate with each other. This could result from one of the controllers being rebooted while the error log was being generated. However, it could be a much more serious error that indicates a problem with the Fibre Channel connections.

- **FCP_ARRAY_ERR14** *ARRAY DRIVE FAILURE*
A serious or unrecoverable error has been detected on a physical disk within the DS4000 subsystem. A system engineer might be able to obtain the exact cause from an analysis of the sense data.
- **FCP_ARRAY_ERR15** *CACHE BATTERY LOW/DATA LOSS POSSIBLE*
If a controller card is replaced, it is likely that the cache batteries will be flat. It can take two days for the cache batteries to be fully recharged. During this time errors are logged in the error log. Do not replace the controller.
- **FCP_ARRAY_ERR16** *CACHE BATTERY CHARGE BELOW 87.5%*
If a controller card is replaced, it is likely that the cache batteries will be flat. It can take two days for the cache batteries to be fully recharged. During this time errors are logged in the error log. Do not replace the controller.
- **FCP_ARRAY_ERR17** *WORLDWIDE NAME CHANGED*
A controller has changed worldwide names (most likely either it was replaced without placing it in the reset state first, or the cabling was changed so that a different controller with the same SCSI ID is on the loop).
- **FCP_ARRAY_ERR18** *RESERVATION CONFLICT*
An operation failed because the disk array logical drive (LUN) is reserved by another host.
- **FCP_ARRAY_ERR19** *SNAPSHOT VOLUME'S REPOSITORY FULL*
The repository capacity limit has been reached. To resolve this error you can increase the repository capacity.
- **FCP_ARRAY_ERR20** *SNAPSHOT OPERATION STOPPED BY ADMIN*
The FlashCopy (snapshot) operation has been disabled or stopped. To resolve this error you can recreate the FlashCopy.
- **FCP_ARRAY_ERR21** *SNAPSHOT REPOSITORY METADATA ERROR*
There was a problem with the metadata of the FlashCopy (snapshot) repository during the FlashCopy operation. To resolve this error you can recreate the FlashCopy.
- **FCP_ARRAY_ERR22** *REMOTE VOL MIRRORING: ILLEGAL I/O ORIGIN*
The primary logical drive received I/O from a remote array, or the secondary logical drive received I/O from other than the primary logical drive. To resolve this error you can try the operation again.
- **FCP_ARRAY_ERR23** *SNAPSHOT OPERATION NOT ALLOWED*
The repository capacity limit has been reached, so the FlashCopy (snapshot) operation has failed. To resolve this error you can delete or recreate the FlashCopy.
- **FCP_ARRAY_ERR24** *SNAPSHOT VOLUME'S REPOSITORY FULL*
The repository capacity limit has been reached. To resolve this error you can delete or recreate the FlashCopy (snapshot).
- **FCP_ARRAY_ERR25** *CACHED DATA WILL BE LOST IF CONTROLLER FAILS*
This message is a warning that a disk array logical drive (LUN) is running with write cache enabled and cache mirroring disabled. The warning displays when the LUN is opened, and it displays again every 24 hours until cache mirroring is enabled again.

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

- **FCP_ARRAY_ERR26 LOGICAL VOLUME IS WRITE PROTECTED**

The status of the logical drive is read-only. The probable reason is that it is a secondary logical drive of a FlashCopy, VolumeCopy, or remote mirror pair. Check which relationship applies to the logical drive.

- For FlashCopy, a status of read-only on the secondary logical drive usually indicates that the repository is full.
- For VolumeCopy, both the primary and secondary logical drives are read-only during the copy. The secondary logical drive is read-only when the copy is stopped but the copy pair had not been deleted.
- For remote mirroring, the secondary logical drive is always read-only, as long as the mirror is active.

Redistributing logical drives in case of failure

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

Complete the following steps to manually redistribute logical drives to their paths:

1. Repair or replace any faulty components. For more information, see the *IBM TotalStorage DS4000 Installation and Support Guide* for the appropriate DS4000 storage server.
2. Using the Subsystem Management window, redistribute logical drives to their preferred paths by clicking **Advanced** → **Recovery** → **Redistribute Logical Drives**.

Note: If a large number of LUNs is configured on the DS4000 system, redistributing logical drives might take 60 minutes or more to complete, depending on how busy the system is.

3. Run the **fget_config** command to verify the active paths, as shown in this example:

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

Chapter 3. Installing and configuring storage management software on HP-UX host systems

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

- “Hardware, firmware, and host system requirements”
- “Installing DS4000 Storage Manager” on page 66
- “Configuring storage subsystems” on page 69
- “Creating direct-attached and SAN-attached configurations” on page 75
- “Enabling multipath I/O with PV-links” on page 75
- “Redistributing logical drives in case of failure” on page 81

Hardware, firmware, and host system requirements

This section lists the the following requirements:

- Hardware and firmware requirements for the DS4000 configuration
- Hardware and software requirements for the HP-UX host system

Note: For the latest installation requirements and user information about Storage Manager, the HP-UX host operating system, and hardware components, see the readme file that is located in the HP-UX DS4000 Installation CD. The most recent copy of the readme file, which supersedes this document, is maintained at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for HP-UX and click the link in the Current Version column.

DS4000 hardware and firmware requirements

Table 17 on page 64 lists the DS4000 Fibre Channel storage servers that you can use with DS4000 Storage Manager.

Table 17. DS4000 storage servers that you can use with HP-UX systems

Product Name	Model	Minimum appware/snapware (firmware) and NVSRAM versions
IBM TotalStorage DS4100 SATA Storage Server	1724-100	For the most recent readme files, see the following Web site: www-1.ibm.com/servers/storage/support/disk/ Click the link for your storage server. When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for HP-UX and click the link in the Current Version column.
IBM TotalStorage DS4500 Fibre Channel Storage Server	1742-90U	
IBM TotalStorage DS4400 Fibre Channel Storage Server	1742-1RU	
IBM TotalStorage DS4300 Standard and Turbo Fibre Channel Storage Servers	1722-60U, 1722-60X	

Table 18 lists the DS4000 EXP storage expansion enclosures and other hardware that you can use with DS4000 Storage Manager.

Note: For important information about using the DS4000 EXP100 SATA storage expansion enclosure, see *IBM TotalStorage DS4000 EXP100 storage expansion enclosure Installation and User's Guide (GC26-7601)*.

Table 18. Supported versions of hardware for HP-UX systems

Hardware	Product Name	Model	Minimum firmware version
storage expansion enclosures	IBM TotalStorage DS4000 EXP100 SATA storage expansion enclosure	1710-10U or 10X	9554
	IBM TotalStorage DS4000 EXP710 Fibre Channel storage expansion enclosure		
	IBM TotalStorage DS4000 EXP700 Fibre Channel storage expansion enclosure	1740-1RU or 1RX	9326
	IBM TotalStorage DS4000 EXP500 Fibre Channel storage expansion enclosure	3560-1RU or 1RX	9166
Host Bus Adapters	HP FC Adapter	A5158A, A6795A	

Table 18. Supported versions of hardware for HP-UX systems (continued)

Hardware	Product Name	Model	Minimum firmware version
Switches	Brocade switch	2109-S08 2109-S16	2.6.2a
	Brocade switch	2109-F16	3.1.2a
	Brocade switch	2109-M12	4.0.2c
	Brocade switch	2109-F32	4.2.0c
	Brocade switch	2109-F08	3.1.2a
	McData switch	2032-064, 2032-140	6.01.00 - 18
	McData switch	2031-016	6.01.00 - 18
	McData switch	2031-216, 2031-224	6.01.00 - 18
	InRange switch	2042-001	4.3.0
	Cisco switch	2062-D01, 2062-D07, 2062-T07	1.3.4a

HP-UX system requirements

This section lists the minimum hardware and software requirements for an HP-UX host in a DS4300, DS4400, or DS4500 storage subsystem that is managed by DS4000 Storage Manager.

Hardware requirements

The HP-UX system must be an HP 9000/Series 800 server with:

- 440 MHz processor or faster
- 1 GB memory
- 0.7 MB available on /opt and root (or root-equivalent) privileges for installing SUtil and SAgent
- 47.5 MB available disk space on /opt and at least 95 MB available on /tmp for installation in an SMruntime environment

Software requirements

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

- HP-UX 11.0 (64 bit) with the following patches:
 - PHKL_23939
 - PHKL_24027
 - PHKL_27364, version 1.0
 - PHKL_29834
 - To run HP-UX SDK for Java 2 Platform applications and applets using GUIs, you must also ensure that the HP C++ runtime libraries are installed on the HP-UX system; the latest version is available as patch PHSS_1658
- HP-UX 11.i (64 bit) with the following patches:
 - PHKL_26743, version 1.0

- PHKL_29365, version 1.0
- PHCO_27408, version 1.0
- For high-availability clusters of HP 9000/Series 800 computers, install the HP MC/Service Guard software package.

Ensure that the following maximum kernel parameters are configured, as shown in Table 19.

Table 19. 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

Installing DS4000 Storage Manager

Use the following procedures, in the order listed, to install the DS4000 Storage Manager software on an HP-UX storage management station or on an HP-UX host acting as a storage management station.

1. “Installing the client software”
2. “Installing the host software” on page 68

Installing the client software

Install the client software in the following order:

1. DS4000 SMruntime software
2. DS4000 SMclient software

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

Prerequisites

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

- This is the HP 9000-series server that you identified as the storage management station.
- This machine meets the minimum hardware and software requirements described in “HP-UX system requirements” on page 65.
- Neither the SMruntime software nor the SMclient software is installed on the host and you are ready to install them on this machine.
- The DS4000 Storage Manager installation files are available in a directory on the system.

Steps for installing SMruntime software

Modify the commands as needed for your installation. No restart is required during the SMruntime installation process.

1. Insert the HP-UX installation CD in the CD-ROM drive.

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

```
# mount -o cdcase -F cdfs /path_name /cdrom
# swinstall -s /cdrom/HP-UX/SMruntime_0901500
```

where *path_name* is the directory path to the SMruntime package.
The Software Selection window opens and displays this message:

```
SMruntime
```

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

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

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

```
# swverify -v SMruntime
```

If no failure is reported, you are done installing SMruntime. If a failure is reported, follow the instructions in the `/var/adm/sw/swagent.log` file.

Result: The SMruntime installation of this HP-UX storage management station is complete.

Steps for installing SMclient software

Modify the commands as needed for your 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-0901500
```

The Software Selection window opens and displays the following message:

```
SMclient
```

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

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

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

```
# swverify -v SMclient
```

If no failure is reported, you are done installing SMclient. If a failure is reported, follow the instructions in the /var/adm/sw/swutil.log file.

Result: The SMclient installation of this HP-UX storage management station is complete.

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

Installing the host software

If you are installing the optional SMagent, install the host software in the following order. If you are not installing SMagent, then SMutil is the only DS4000 Storage Manager host software that you need to install.

- DS4000 SMagent software (optional - for in-band management only)
- DS4000 SMutil software (required)

Note: SMagent and SMutil are dependent on SMruntime, which is a Java compiler for the DS4000 Storage Manager software and must be installed first. If you do not install SMruntime you receive errors during the software install analysis phase.

Prerequisites

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

- This is the HP 9000-series server that you identified as the storage management station.
- This machine meets the minimum hardware and software requirements described in “HP-UX system requirements” on page 65.
- SMruntime has been installed.

Note: SMagent is not supported with Storage Manager controller firmware version 5.42.xx.xx.

Steps for installing SMagent software (optional)

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

Modify the commands as needed for your installation. No restart is required during the SMagent installation process.

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

```
# swinstall -s /cdrom/HP-UX/SMagent-HP-0901500
```

The Software Selection window opens and displays this message:

```
SMagent
```

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

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

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

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

```
# swverify -v SMagent
```

If no failure is reported, you are done installing SMagent. If a failure is reported, follow the instructions in the `/var/adm/sw/swagent.log` file.

Result: The SMagent installation on this HP-UX storage management station is complete.

Steps for installing SMutil software (required)

Modify the commands as needed for your installation. No restart is required during the SMutil installation process.

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

```
# swinstall -s /cdrom/HP-UX/SMutil_0901500
```

The Software Selection window opens and displays this message:

```
SMutil
```

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

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

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

```
# swverify -v SMutil
```

If no failure is reported, you are done installing SMutil. If a failure is reported, follow the instructions in the `/var/adm/sw/swutil.log` file.

Result: The SMutil installation of this HP-UX storage management station is complete.

Configuring storage subsystems

After you install DS4000 Storage Manager, you need to configure the subsystem. System configuration includes configuring logical drives and creating storage arrays using the Subsystem Management window of DS4000 Storage Manager.

See “Configuring storage subsystems” on page 139 for additional information about storage configuration.

Before you begin: Note the following information:

- All DS4000 storage servers have an access logical drive set to LUN 31.

Note: Do not delete the access LUN or the access logical drive.

- AVT/ADT is enabled by default; leave it in that state for HP-UX hosts.
- Versions of firmware and NVSRAM that come with this product might be later releases than those described in this document. To ensure that you have the latest versions of the firmware and the NVSRAM, read the readme file that is provided with the product. You can also find the latest readme file at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for HP-UX and click the link in the Current Version column.

Note: If you do not have the latest versions of firmware and NVSRAM, download them and complete the procedure described in “Upgrading controller firmware and NVSRAM” on page 71. If the version numbers are current, you can skip that procedure.

- If you have purchased a Storage Partitioning premium feature, ensure that the feature is enabled. See “Enabling your premium features” on page 21 for more information.

Complete the following procedures to configure DS4000 Storage Manager for an HP-UX system.

Adding storage subsystems to SMclient

Complete the following steps to add storage subsystems to SMclient:

1. To set up the storage subsystem for HP-UX, the subsystem must be physically configured, at least initially, for direct management through the Ethernet connections on each controller as well as through the Fibre Channel connection. Install the DS4000 Storage Manager software (SMruntime, SMclient, SMagent, and SMutil) before configuring the subsystem.
2. After the disk subsystem is configured on the network, type the following command to run SMclient on either the host server, if it is on the same network as the storage subsystem, or on another machine in the network that has the DS4000 Storage Manager software installed:

```
# SMclient
```

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

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

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

Upgrading controller firmware and NVSRAM

The procedures in this section describe how to download and install DS4000 controller firmware and NVSRAM.

Checking current versions of firmware and NVSRAM

Versions of firmware and NVSRAM that are provided with this product might be later releases than those described in this document.

To check the versions of firmware and NVSRAM that are currently installed on your system, complete the following steps:

1. In the Subsystem Management window, click **Storage Subsystem** → **View Profile**. The Storage Subsystem Profile window displays.
2. In the Storage Subsystem Profile window, click the Controllers tab to see the current versions of firmware and NVSRAM.
3. To find whether the versions of firmware and NVSRAM that are currently on your storage subsystem are the latest versions, see the readme file at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for HP-UX and click the link in the Current Version column.

The latest versions are listed in the readme file.

4. If you find that you do not have the latest versions of the firmware and NVSRAM, click the **Back** button on your browser to return to the previous page. From there, you can download the correct versions to the host by clicking the appropriate link that is listed in the IBM DS4000 Controller Firmware matrix.

Upgrading firmware using traditional or staged download

There are two methods for downloading the firmware to the controllers:

- “Traditional download”
- “Staged download”

You can select which method you want to use during the upgrade process, as described in “Steps for upgrading the controller firmware” on page 72.

Traditional download: The traditional download process takes significantly longer and must be done in one phase, rather than in two phases as with the staged controller firmware download. Therefore the preferred method of controller firmware download is the staged controller firmware download.

Staged download: With staged firmware download, the download process occurs in two distinct phases: firmware transfer and firmware activation. You can perform the time-consuming task of loading the firmware online so that it is functionally transparent to the application. You can then defer the activation of the loaded firmware to a convenient time.

Some reasons for activating firmware files at a later time include the following:

- **Time of day** - Activating can take a long time, so you might want to wait until I/O loads are lighter because the controllers will go offline briefly to activate the new firmware.

- **Type of package** - You might want to test the new firmware package on one storage subsystem before activating it on other storage subsystems.
- **Multiple controllers** - You can download controller firmware from the storage management software to all storage subsystem controllers for later activation.

Notes:

1. You cannot use staged download to upgrade Storage Manager controller firmware version 5.4x.xx.xx to version 6.1x.xx.xx. You can only use staged download if the Storage Manager controller firmware version is already 6.1x.xx.xx, and you are downloading another version of 6.1x.xx.xx firmware.
2. Staged controller firmware download is not supported on DS4400 storage servers.

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

Steps for upgrading the controller firmware: Complete the following steps to install a new firmware version on the controllers:

Before you begin:

1. Download the correct version of the firmware from the Web site that is listed in “Checking current versions of firmware and NVSRAM” on page 71. Place the file in a designated directory on the host system.
2. Ensure that you have management connections to both controllers, and that controllers are in an optimal state.
 1. Stop I/O to the controllers.
 2. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.
3. In the Subsystem Management window, click **Advanced** —> **Maintenance** —> **Download** —> **Controller Firmware**. The Download Firmware window displays.
4. Click **Browse** to select the firmware file that you previously downloaded to a directory on the host.
5. If you want to perform a staged firmware download, select the check box next to “Transfer files but don’t activate them (activate later).” If you want to perform a traditional firmware download, do not select the check box.
6. Click **OK**. The Confirm Download window displays.
7. Click **Yes** to start the download.
8. Close the Subsystem Management window and then reopen it to complete the firmware update.
9. Type the following command to restart the host:

```
# reboot
```

10. Verify that the firmware was successfully installed.
11. Restart I/O.

Upgrading the NVSRAM

Complete the following steps to install a new version of NVSRAM on the controllers. The NVSRAM will be activated immediately.

Before you begin:

1. Download the correct NVSRAM version from the Web site that is listed in “Checking current versions of firmware and NVSRAM” on page 71. Place the file in a designated directory on the host system.
2. Ensure that you have management connections to both controllers, and that controllers are in an optimal state.
 1. Stop I/O to the controllers.
 2. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.
3. In the Subsystem Management window, click **Advanced** → **Maintenance** → **Download** → **Controller NVSRAM**. The Download NVSRAM window displays.
4. Click **Browse** to select the NVSRAM file that you previously downloaded to a directory on the host.
5. Click **OK**. The Confirm Download window displays.
6. Click **Yes** to start the download.
7. Verify that the NVSRAM was successfully installed.

Storage Partitioning: Defining an HP-UX host group

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

Notes:

1. Create the HP-UX host group at the storage subsystem level. Do not create host groups at the default group level.

Exception: If you are running a DS4300 configuration without partitioning enabled, you can use the default host group.
2. In a cluster partition, perform logical drive mappings on the host group level so that all the hosts can see the same storage. In a normal partition, perform logical drive mappings on the host level.

For more information about Storage Partitioning, see “Storage Partitioning premium feature” on page 21.

Before you begin: Identify the worldwide port name for each of your HBAs. Look for the worldwide port name for the HBA and write down the number for future reference. Use the following command to find your HBAs listed in the dev directory as td0, td1, and so on, where X in the following example represents the td number.

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

Steps for defining a host group

Complete the following steps to create an HP-UX host group:

1. Click the **Mappings View** tab on the Subsystem Management window.
2. In the Mappings window, create a new host group by clicking **Mappings** → **Define** → **Host Group**.

Note: Make sure that the storage subsystem is highlighted in the left panel of the Subsystem Management window. Do not click on Undefined Mappings.

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, then right-click and select **Define Host Port**.

Note: You might be required to restart the host if you cannot see the WWPN for your HBAs. Restarting causes the system to log into the switch. If the restart does not work, manually enter the WWPNs when prompted.

- c. Type or select the desired host port for the first HBA, and then change the host type to **HP-UX**. Click **Add**.

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

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

Mapping LUNs to a storage partition

Before you begin: Read “DS4000 Storage Manager performance planning” on page 18 for information about planning the number of LUNs to map to each partition.

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

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

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

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

Note: It is recommended that you do not map at the host group level. If you do, every host can see all the LUNs.

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

Creating direct-attached and SAN-attached configurations

DS4000 Storage Manager supports IBM DS4000 Fibre Channel Storage Servers in direct-attached HP-UX configurations or in a SAN environment through switches in HP-UX configurations.

Creating a direct-attached configuration

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

- One or two HP-UX servers can be connected to DS4000 storage servers.
- Two server DS4400 or DS4500 configurations require four host-side minihubs, each with exactly one Fibre Channel connection from each HBA to a minihub.

Note: DS4300 Turbo storage servers do not have minihubs.

- There are two or four HBAs (HP A6795A or HP A5158A) per DS4000 storage server. Each pair must be configured to one DS4000 partition.
- No external hubs are being used.

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

1. Connect the HBAs to each controller or minihub port of the DS4000 storage server.
2. Configure and verify the configuration.

Creating a SAN-attached configuration

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

- Multiple Fibre Channel HBAs within the same server cannot “see” the same DS4000 controller port.
- The IBM Fibre Channel HBAs are isolated from each other if they are connected to the same switch that is connected to the same DS4000 controller port.
- Each Fibre Channel HBA and controller port must be in its own fabric zone, if they are connecting through a single Fibre Channel switch, such as a 2109-F16.

Multiple DS4000 devices can be configured to the same set of Fibre Channel HBAs through a Fibre Channel switch.

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

Note: For more information about zoning and enabling zones, see Appendix E, “Connecting HBAs in a FC switch environment,” on page 165.

1. Connect the HBAs to the switch or switches.
2. Connect the DS4000 storage subsystems to the switch or switches.
3. Set the required zones on the Fibre Channel switch or switches, if applicable.
4. Configure and verify the configuration.

Enabling multipath I/O with PV-links

If the HP-UX system is attached with two host bus adapters to the DS4000 storage server, you can establish redundant access to storage by using physical volume links (PV-links), a feature of the HP-UX operating system. PV-links achieve access redundancy by using devices with both primary and secondary paths to the same device.

Important:

- There are two methods for establishing redundant access to storage using PV-links:
 - If you have Storage Manager controller firmware version 6.1.xx.xx or 5.xx.xx.xx installed, use method 1.
 - If you have Storage Manager controller firmware version 4.xx.xx.xx installed, use method 2.
- For both methods, you must have SMutil installed on the host.

Using PV-links: Method 1

If you have Storage Manager controller firmware version 6.1.xx.xx or 5.xx.xx.xx installed, use the following procedure to enable multipath I/O by using PV-links:

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

```
#hot_add
```

2. Run the **SMdevices** command. The system provides a dump similar to the example in Figure 11 on page 77. Notice that every logical drive and logical drive access unit has been assigned a logical unit number (LUN). Each logical drive has two paths. Each RAID controller has one logical drive access. For example, a subsystem that contains two RAID controllers has two logical drive accesses.

```
#SMdevices
```



```

/dev/rdisk/c166t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>,
Preferred Path (Controller-B): In Use]
/dev/rdisk/c166t0d1 [Storage Subsystem DS4000, Logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>,
Alternate Path (Controller-B): Not In Use]
/dev/rdisk/c166t0d2 [Storage Subsystem DS4000, Logical Drive Finance,
LUN 2, Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>,
Alternate Path (Controller-B): Not In Use]
/dev/rdisk/c166t0d3 [Storage Subsystem DS4000, Logical Drive Purchasing,
LUN 3, Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>,
Alternate Path (Controller-B): Not In Use]
/dev/rdisk/c166t0d4 [Storage Subsystem DS4000, Logical Drive Development,
LUN 4, Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>,
Preferred Path (Controller-B): In Use]
/dev/rdisk/c166t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f56d00000001b00000000>]

/dev/rdisk/c172t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>,
Alternate Path (Controller-A): Not In Use]
/dev/rdisk/c172t0d1 [Storage Subsystem DS4000, logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>,
Preferred Path (Controller-A): In Use]
/dev/rdisk/c172t0d2 [Storage Subsystem DS4000, Logical Drive Finance, LUN 2,
Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>,
Preferred Path (Controller-A): In Use]
/dev/rdisk/c172t0d3 [Storage Subsystem DS4000, Logical Drive Purchasing, LUN 3,
Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>,
Preferred Path (Controller-A): In Use]
/dev/rdisk/c172t0d4 [Storage Subsystem DS4000, Logical Drive Development, LUN 4,
Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>,
Alternate Path (Controller-A): Not In Use]
/dev/rdisk/c172t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f5d6c0000002200000000>]

```

Figure 11. Sample `SMdevices` command output (method 1)

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

```
#reboot
```

- Determine the preferred and alternate path for each logical drive by examining the output from the `SMdevices` command, as shown in the example in Figure 11. Notice that each device is listed twice; one instance is the preferred path and one instance is the alternate path.

Preferred path

In the sample output that is shown in Figure 11, the preferred path is `/dev/rdsk/c166t0d0`:

```

/dev/rdsk/c166t0d0 [Storage Subsystem DS4000, Logical Drive
Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>,
Preferred Path (Controller-B): In Use]

```

Alternate path

In the sample output that is shown in Figure 11, the alternate path is `/dev/rdsk/c172t0d0`:

```
/dev/rdisk/c172t0d0 [Storage Subsystem DS4000, Logical Drive
Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>,
Alternate Path (Controller-A): NotIn Use]
```

Using PV-links: Method 2

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

- Determine the preferred and alternate paths
- Create the logical drives and logical drive groups

Determining preferred and alternate paths

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

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

```
#hot_add
```

2. Run the **SMdevices** command. The system provides a dump similar to the example in Figure 12. Notice that every logical drive and logical drive access unit has been assigned a logical unit number (LUN). Each logical drive has two paths. Each RAID controller has one logical drive access. For example, a subsystem that contains two RAID controllers has two logical drive accesses.

```
#SMdevices
```

```
/dev/rdisk/c166t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>]
/dev/rdisk/c166t0d1 [Storage Subsystem DS4000, Logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>]
/dev/rdisk/c166t0d2 [Storage Subsystem DS4000, Logical Drive Finance, LUN 2,
Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>]
/dev/rdisk/c166t0d3 [Storage Subsystem DS4000, Logical Drive Purchasing, LUN 3,
Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>]
/dev/rdisk/c166t0d4 [Storage Subsystem DS4000, Logical Drive Development, LUN 4,
Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>]
/dev/rdisk/c166t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f56d00000001b0000000>]

/dev/rdisk/c172t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>]
/dev/rdisk/c172t0d1 [Storage Subsystem DS4000, logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>]
/dev/rdisk/c172t0d2 [Storage Subsystem DS4000, Logical Drive Finance, LUN 2,
Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>]
/dev/rdisk/c172t0d3 [Storage Subsystem DS4000, Logical Drive Purchasing, LUN 3,
Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>]
/dev/rdisk/c172t0d4 [Storage Subsystem DS4000, Logical Drive Development, LUN 4,
Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>]
/dev/rdisk/c172t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f5d6c000000220000000>]
```

Figure 12. Sample SMdevices command output (method 2)

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

```
#reboot
```

- Determine the preferred and alternate path for each logical drive by examining the output from the **SMdevices** command, as shown in the example in Figure 12 on page 78.

Notice that each device is listed twice; one instance is the preferred path and one instance is the alternate path. Also, notice that each device has a worldwide name (WWN). Part of the WWN of each logical drive is unique for each controller in the DS4000 storage server. If you examine the WWNs for the logical drive access in Figure 12 on page 78, you notice that they differ in only five digits, f56d0 and f5d6c.

The devices in Figure 12 on page 78 are viewed through the controllers c166 and c172. To determine the preferred path of a specific logical drive seen by the operating system perform the following steps:

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

```
/dev/rdsk/c166t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31, Logical Drive WWN <600a0b80000f56d0000001b00000000>]
```

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

```
/dev/rdsk/c172t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31, Logical Drive WWN <600a0b80000f5d6c0000002200000000>]
```

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

```
/dev/rdsk/c166t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0, Logical Drive g<600a0b80000f56d00000001e3eaead2b>]
```

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

```
/dev/rdsk/c172t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>]
```

- To keep a record for future reference, enter this path information for LUN 0 into a matrix (similar to the one in Table 20).

Table 20. Sample record of logical drive preferred and alternate paths

LUN	Logical drive name	Preferred path	Alternate path
0	Accounting	/dev/rdsk/c166t0d0	/dev/rdsk/c172t0d0
1	HR	/dev/rdsk/c172t0d1	/dev/rdsk/c166t0d1
2	Finance	dev/rdsk/c172t0d2	/dev/rdsk/c166t0d2
3	Purchasing	/dev/rdsk/c172t0d3	/dev/rdsk/c166t0d3
4	Development	/dev/rdsk/c166t0d4	/dev/rdsk/c172t0d4

- d. Repeat step 3a on page 79 through step 3c on page 79 for each logical drive that is seen by the operating system.

Continue to the next procedure.

Creating volumes and volume groups

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

Recommendation: Do not use SAM for DS4000 storage configuration. If you do, you might get unexpected results.

Note: The steps in this procedure refer to LUN 0 in Table 20 on page 79.

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

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

The system confirms the creation of the new physical volume.

2. Create volume groups.

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

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

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

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

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

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

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

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

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

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

- 1) Add the primary path for LUN1.

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

2) Add the secondary path for LUN1.

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

3. Create logical volumes. For more information, refer to HP-UX documentation.
4. Create file systems for the logical volumes.
5. Repeat step 1 on page 80 through step 4 to create additional volume groups. For more information, refer to HP-UX documentation.
6. Verify the primary (preferred) and secondary (alternate) paths for each device by typing the following command:

```
#vgdisplay -v vgname
```

where *vgname* is the volume group name.

Redistributing logical drives in case of failure

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

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

1. Repair or replace any faulty components. For more information, see the *IBM TotalStorage DS4000 Installation and Support Guide* for the appropriate DS4000 storage server.
2. Using the Subsystem Management window, redistribute logical drives to their preferred paths by clicking **Advanced** → **Recovery** → **Redistribute Logical Drives**.

Chapter 4. Installing and configuring storage management software on Solaris host systems

This chapter provides the following information for Solaris operating systems:

- “Hardware, firmware, and host system requirements”
- “Installing DS4000 Storage Manager” on page 86
- “Configuring storage subsystems” on page 90
- “Creating direct-attached and SAN-attached configurations” on page 95
- “Configuring the host bus adapter cards” on page 96
- “Installing a failover driver on Solaris” on page 105
- “Verifying external storage” on page 118
- “Redistributing logical drives in case of failure” on page 118

Hardware, firmware, and host system requirements

This section lists the the following requirements:

- Hardware and firmware requirements for the DS4000 configuration
- Hardware and software requirements for the Solaris host system

Note: For the latest installation requirements and user information about Storage Manager, the Solaris host operating system, and hardware components, see the readme file that is located in the Solaris DS4000 Installation CD. The most recent copy of the readme file, which supersedes this document, is maintained at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for Solaris and click the link in the Current Version column.

DS4000 hardware and firmware requirements

Table 21 on page 84 lists the DS4000 Fibre Channel storage servers that you can use with DS4000 Storage Manager.

Table 21. DS4000 storage servers that you can use with Solaris systems

Product name	Model	Minimum appware/snapware (firmware) and NVSRAM versions
IBM TotalStorageDS4100 SATA Storage Server	1724-100	For the most recent readme files, see the following Web site: www-1.ibm.com/servers/storage/support/disk/ Click the link for your storage server. When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for Solaris and click the link in the Current Version column.
IBM TotalStorage DS4500 Fibre Channel Storage Server	1742-90U	
IBM TotalStorage DS4400 Fibre Channel Storage Server	1742-1RU	
IBM TotalStorage DS4300 Standard and Turbo Fibre Channel Storage Servers	1722-60U, 1722-60X	

Table 22 lists the DS4000 EXP storage expansion enclosures and other hardware that you can use with DS4000 Storage Manager.

Note: For important information about using the DS4000 EXP100 SATA storage expansion enclosure, see *IBM TotalStorage DS4000 EXP100 storage expansion enclosure Installation and User's Guide (GC26-7601)*.

Table 22. Hardware you can use with Solaris systems

Hardware	Product Name	Model	Minimum firmware version
storage expansion enclosures	IBM TotalStorage DS4000 EXP100SATA storage expansion enclosure	1710-10U or 10X	9554
	IBM TotalStorage DS4000 EXP710Fibre Channel storage expansion enclosure		
	IBM TotalStorage DS4000 EXP700 Fibre Channel storage expansion enclosure	1740-1RU or 1RX	9326
	IBM TotalStorage DS4000 EXP500 Fibre Channel storage expansion enclosure	3560-1RU or 1RX	9166

Table 22. Hardware you can use with Solaris systems (continued)

Hardware	Product Name	Model	Minimum firmware version
Host Bus Adapters	JNI FC Adapter	FC64-1063	
	JNI FC Adapter	FCE-1063 S-bus 64 bit	
	JNI FC Adapter	FCE-1473 S-bus 64 bit	3.8.9
	JNI FC Adapter	FCE2-1473 S-bus 64 bit	3.8.9
	JNI FC Adapter	FCE-6410-N 33 Mhz	
	JNI FC Adapter	FCE-6460-PCI 64 bit	3.8.9
	JNI FC Adapter	FCE2-1063 S-bus 64 bit	
	JNI FC Adapter	FCE2-6412 66 Mhz 64 bit	
	JNI FC Adapter	FCI-1063	
	JNI FC Adapter	FCC-6562	3.8.9
	JNI FC Adapter	FCC2-6562	3.8.9
	JNI FC Adapter	FCX-6562	3.8.9
	JNI FC Adapter	FCC\X2-6562	3.8.9
	QLogic Adapter	4.08	2.00.01
Switches	Brocade switch	2109-S08, 2109-S16	2.6.2a
	Brocade switch	2109-F16	3.1.2a
	Brocade switch	2109-F8	3.1.2a
	Brocade switch	2109-M12	4.0.2c
	Brocade switch	2109-F32	4.2.0c
	McData switch	2032-064, 2032-140	6.01.00 - 18
	McData switch	2031-016	6.01.00 - 18
	McData switch	2031-216, 2031-224	6.01.00 - 18
	InRange switch	2042-001	4.3.0
	Cisco switch	2062-D01, 2062-D07, 2062-T07	1.3.4a

Solaris system requirements

This section lists the minimum hardware and software requirements for a Solaris host in a DS4100, DS4300, DS4400, or DS4500 storage subsystem that is managed by DS4000 Storage Manager.

Hardware requirements

The Solaris system must be a Sparc S20 processor with:

- 1 GB system memory
- CD-ROM drive
- Mouse or similar pointing device
- Ethernet network interface card
- 1 MB available on /opt and root (or root-equivalent) privileges for installing RDAC

Software requirements

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

- Solaris 8 with the following patches (minimum versions):
 - Solaris 8 Recommended patch cluster
- Solaris 9 with the following patches (minimum versions):
 - 113454-14
- 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.

Installing DS4000 Storage Manager

Use the following procedures, in the order listed, to install the DS4000 Storage Manager software on a Solaris storage management station.

1. “Installing the client software”
2. “Installing the host software” on page 88

Installing the client software

Install the client software in the following order:

1. DS4000 SMruntime software
2. DS4000 SMclient software

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

Prerequisites

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

- This is the SUN SPARCstation workstation that you have identified as the storage management station.
- This machine is running Solaris 8 or Solaris 9, and meets the minimum hardware and software requirements described in “Solaris system requirements” on page 85.
- Neither DS4000 SMruntime software nor DS4000 SMclient software is installed on the host, and you are ready to install them on this machine.
- The DS4000 Storage Manager installation files are available in a directory on the system.

Steps for installing SMruntime software

Modify these commands as needed for your installation. No restart is required during the client installation process.

Before you begin: Visit the Web site that is listed in “Hardware, firmware, and host system requirements” on page 83 to ensure you have the most recent version of the software.

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

```
pkgadd -d path/filename.pkg
```

where *path/filename* is the directory path and name of the package that you want to install. Information about packages that can be installed in the specified directory is displayed on the command line. The following is an example of what you might see displayed.

```
The following packages are available:  
1 SMruntime  
IBM DS4000 Storage Manager 9 Runtime (sparc)  
version numberSelect package(s) you wish to process (or 'all'  
to process all Packages). (default:all) [?,?,q]:
```

3. Press **Enter** to select the default (all). The installation process begins. The following prompt is displayed:

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

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

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

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

```
# pkginfo -l SMruntime
```

The installation is successful if the following message displays:

```
INSTDATE: date/time  
STATUS: completely installed
```

where *date/time* is today’s installation date and time.

Steps for installing SMclient software

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

Modify these commands as needed for your installation. 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 path/filename.pkg
```

where *path/filename* is the directory path and name of the package that you want to install.

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

```
The following packages are available:

1 SMclient                IBM DS4000 Storage Manager 9 Client
                           (sparc) version number

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

3. Press **Enter** to select the default (all).

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 is successful:

```
# pkginfo -l SMclient
```

The installation is successful if the following message displays:

```
INSTDATE: date/time
STATUS: completely installed
```

where *date/time* is today's installation date and time.

Result: The client software installation on this Solaris storage management station is complete. Remove the installation CD from the CD-ROM drive.

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

Installing the host software

If you are installing the optional SMagent, install the host software in the following order. If you are not installing SMagent, then SMutil is the only DS4000 Storage Manager host software that you need to install.

- DS4000 SMagent software (optional - for in-band management only)
- DS4000 SMutil software (required)

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 8 or Solaris 9, and meets the minimum hardware and software requirements described in “Solaris system requirements” on page 85.
- The SMagent, SMutil, and RDAC software packages are not installed, and you are ready to install them on this machine.

Note: SMagent is not supported with Storage Manager controller firmware version 5.42.xx.xx.

Steps for installing SMagent software (optional)

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

Complete the following steps to install SMagent:

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

```
pkgadd -d path/filename.pkg
```

where *path/filename* is the directory path and name of the package that you want to install.

The installation process begins.

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

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

```
1 SMagent                IBM DS4000 Storage Manager 9 Agent
                        (sparc) version number
```

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

2. Press **Enter** to select the default (all).

The installation process begins. The following prompt is displayed:

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

3. Type **y** and press **Enter**.

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

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

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

```
# pkginfo -l SMagent
```

The installation is successful if the following message displays:

```
INSTDATE: date/time
STATUS: completely installed
```

where *date/time* is today's installation date and time.

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

Steps for installing SMutil software (required)

Complete the following steps to install SMutil:

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

```
pkgadd -d path/filename.pkg
```

where *path/filename* is the directory path and name of the package that you want to install.

The installation process begins.

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

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

2. Type the following command to verify that the installation is successful:

```
# pkginfo -l SMutil
```

The installation is successful if the following message displays:

```
INSTDATE: date/time  
STATUS: completely installed
```

where *date/time* is today's installation date and time.

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

Configuring storage subsystems

To configure DS4000 Storage Manager for a Solaris system, complete the following procedures in the order that they are described in this section:

- Add storage subsystems to DS4000 SMclient
- Update NVSRAM and DS4000 firmware (if necessary)

Before you begin: Read the following information:

- All DS4000 storage servers have an access logical drive set to LUN 31.

Note: Do not delete the access LUN or the access logical drive.

- If you have purchased a Storage Partitioning premium feature, ensure that the feature is enabled. See "Enabling your premium features" on page 21 for more information.
- Versions of firmware and NVSRAM that are provided with this product might be later releases than those described in this document. To ensure that you have the latest versions of the firmware and the NVSRAM, read the readme file that is provided with the product. You can also find the latest readme file at the Web site that is appropriate for your DS4000 storage server:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for Solaris and click the link in the Current Version column.

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

Adding storage subsystems to SMclient

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

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

Upgrading controller firmware and NVSRAM

The procedures in this section describe how to download and install DS4000 controller firmware and NVSRAM.

Checking current versions of firmware and NVSRAM

Versions of firmware and NVSRAM that are provided with this product might be later releases than those described in this document.

To check the versions of firmware and NVSRAM that are currently installed on your system, complete the following steps:

1. In the Subsystem Management window, click **Storage Subsystem** —> **View Profile**. The Storage Subsystem Profile window displays.
2. In the Storage Subsystem Profile window, click the Controllers tab to see the current versions of firmware and NVSRAM.
3. Click **Close** to close the Storage Subsystem Profile window.
4. To find whether the versions of firmware and NVSRAM that are currently on your storage subsystem are the latest versions, see the readme file at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for Solaris and click the link in the Current Version column.

The latest versions are listed in the readme file.

5. If you find that you do not have the latest versions of the firmware and NVSRAM, click the **Back** button on your browser to return to the previous page. You can download the correct versions to a directory by clicking the appropriate link that is listed in the IBM DS4000 Controller Firmware matrix.

Upgrading firmware using traditional or staged download

There are two methods for downloading the controller firmware:

- “Traditional download” on page 92
- “Staged download” on page 92

You can select which method you want to use during the upgrade process, as described in “Upgrading the controller firmware.”

Traditional download: The traditional download process takes significantly longer and must be done in one phase, rather than in two phases as with the staged controller firmware download. Therefore the preferred method of controller firmware download is the staged controller firmware download.

Staged download: With staged firmware download, the download process occurs in two distinct phases: firmware transfer and firmware activation. You can perform the time-consuming task of loading the firmware online so that it is functionally transparent to the application. You can then defer the activation of the loaded firmware to a convenient time.

Some reasons for activating firmware files at a later time include the following:

- **Time of day** - Activating can take a long time, so you might want to wait until I/O loads are lighter because the controllers will go offline briefly to activate the new firmware.
- **Type of package** - You might want to test the new firmware package on one storage subsystem before activating it on other storage subsystems.
- **Multiple controllers** - You can download controller firmware from the storage management software to all storage subsystem controllers for later activation.

Notes:

1. You cannot use staged download to upgrade Storage Manager controller firmware version 5.4x.xx.xx to version 6.1x.xx.xx. You can only use staged download if the Storage Manager controller firmware version is already 6.1x.xx.xx, and you are downloading another version of 6.1x.xx.xx firmware.
2. Staged controller firmware download is not supported on DS4000 storage servers.

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

Upgrading the controller firmware

Complete the following steps to install a new firmware version on the controllers:

Before you begin:

1. Download the correct version of the firmware from the Web site that is listed in “Checking current versions of firmware and NVSRAM” on page 91. Place the file in a designated directory on the host system.
2. Ensure that you have management connections to both controllers, and that controllers are in an optimal state.
 1. Stop I/O to the controllers.
 2. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.
 3. In the Subsystem Management window, click **Advanced** —> **Maintenance** —> **Download** —> **Controller Firmware**. The Download Firmware window displays.
 4. Click **Browse** to select the firmware file that you previously downloaded to a directory on the host.

5. If you want to perform a staged firmware download, select the check box next to "Transfer files but don't activate them (activate later)." If you want to perform a traditional firmware download, do not select the check box.
6. Click **OK**. The Confirm Download window displays.
7. Click **Yes** to start the download.
8. Close the Subsystem Management window and then reopen it to complete the firmware update.
9. Verify that the firmware was successfully installed.
10. Restart I/O.

Upgrading the NVSRAM

Complete the following steps to install a new NVSRAM version on the controllers. The NVSRAM will be activated immediately.

Before you begin:

1. Download the correct version of the firmware from the Web site that is listed in "Checking current versions of firmware and NVSRAM" on page 91. Place the file in a designated directory on the host system.
2. Ensure that you have management connections to both controllers, and that controllers are in an optimal state.
 1. Stop I/O to the controllers before you upgrade the firmware.
 2. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.
3. In the Subsystem Management window, click **Advanced** → **Maintenance** → **Download** → **Controller NVSRAM**. The Download NVSRAM window displays.
4. Click **Browse** to select the NVSRAM file that you previously downloaded to a directory on the host.
5. Click **OK**. The Confirm Download window displays.
6. Click **Yes** to start the download.
7. Verify that the NVSRAM was successfully installed.

Storage Partitioning: Defining a Solaris host group

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

Notes:

1. Create the Solaris host group at the storage subsystem level. Do not create host groups at the default group level.

Exception: If you are running a DS4300 configuration without partitioning enabled, you can use the default host group.
2. In a cluster partition, perform logical drive mappings on the host group level so that all the hosts can see the same storage. In a normal partition, perform logical drive mappings on the host level.

For more information about storage partitioning, see "Storage Partitioning premium feature" on page 21.

Steps for defining a host group

Complete the following steps to define a Solaris host group:

1. Re-scan the storage subsystem, and then click **Manage the Device** to return to the Subsystem Management window.
2. In the Subsystem Management window, click the **Mappings View** tab.
3. In the Mappings window, create a new host group by clicking **Mappings** → **Define** → **Host Group**.

Note: Make sure that the storage subsystem is highlighted in the left panel of the Subsystem Management window. Do not click on Undefined Mappings.

4. Type the name of the new host group (for example, Solaris). Click **Add**, and then click **Close**.
5. Select the new host group and click **Mappings** → **Define** → **Host**.
6. Define the new host. Type the name of the Solaris host to which the storage subsystem is attached.
 - a. Click **Add**, and then click **Close**.
 - b. Highlight the host that you just added, then right-click and select **Define Host Port**.
 - c. Type or select the desired host port for the first HBA, and then change the host type to **Solaris**. Click **Add**.

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

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

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

- e. Click **File** → **Exit**.

Mapping LUNs to a storage partition

Before you begin: Read “DS4000 Storage Manager performance planning” on page 18 for information about planning the number of LUNs to map to each partition.

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

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

Note: It is recommended that you do not map at the host group level. If you do, every host can see all the LUNs.

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

Creating direct-attached and SAN-attached configurations

DS4000 Storage Manager supports IBM DS4000 Fibre Channel Storage Servers in direct-attached Solaris configurations or in a SAN environment through switches in Solaris configurations.

Creating a direct-attached configuration

In a direct-attached configuration, one or two Solaris servers can be connected to DS4000 storage servers.

Requirements:

- Two-server DS4400 or DS4500 configurations require four host-side minihubs, each with exactly one Fibre Channel connection from each HBA to a minihub.

Note: DS4300 Turbo storage servers do not have minihubs.

- No external hubs can be used.

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

1. Connect the HBAs to each controller or minihub port of the DS4000 storage server.
2. Configure and verify the configuration.

Creating a SAN-attached configuration

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

Requirements:

- Multiple HBAs within the same server must be unable to “see” the same DS4000 controller port.
- The JNI HBAs must be isolated from each other if they are connected to the same switch that is connected to the same DS4000 controller port.
- Each HBA and controller port must be in its own fabric zone, if they are connecting through a single switch, such as a 2109-F16.

See the documentation provided by the switch manufacturer for more information about zoning. Multiple DS4000 devices can be configured to the same set of HBAs through a Fibre Channel switch.

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

+
+

Note: For more information about zoning and enabling zones, see Appendix E, “Connecting HBAs in a FC switch environment,” on page 165.

1. Connect the HBAs to the switch or switches.
2. Connect the DS4000 storage subsystems to the switch or switches.
3. Set the required zones on the Fibre Channel switch or switches, if applicable.

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

4. Configure and verify the configuration.

Configuring the host bus adapter cards

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

JNI

- “Installing the JNI adapter driver package”
- “Modifying the JNI HBA settings” on page 97 (SAN-attached only)

QLogic

- “Installing the QLogic adapter driver package” on page 102
- “Modifying the QLogic HBA settings” on page 104

Note: The QLogic HBA is supported only on Solaris 8 and 9 operating systems.

Installing the JNI adapter driver package

Before you begin:

- Ensure that the JNI HBAs have been installed and attached directly to the controllers or to a SAN fabric switch.
- If attached to SAN fabric switch, ensure that the zones for the DS4000 storage subsystem have been created and enabled.

Notes:

1. All fabric switches must be zoned in such a way that a single HBA can access only one controller per storage array.
2. For more information about zoning and enabling zones, see Appendix E, “Connecting HBAs in a FC switch environment,” on page 165.

The procedures for installing the JNI adapter driver package and modifying the JNI HBA settings both refer to Table 23.

Table 23. JNI model numbers, adapter drivers, and configuration files

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

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

1. In Table 23, identify the driver package for the installed HBAs.
2. Download the most current adapter driver package from the following Web site:

www.jni.com/Drivers

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

```
#pkgadd -d adapter_driver_package
```

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

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

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

```
#pkginfo adapter_driver_package
```

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

```
#pkginfo -l JNIC146x.pkg
```

Modifying the JNI HBA settings

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

This section contains the following procedures:

- “Modifying the JNI HBA settings for McData, Cisco, InRange, and Brocade 1 GB switches”
- “Modifying the JNI HBA settings for a Brocade 2 Gb switch” on page 99

Attention: If you have a direct-attached configuration, skip this section and proceed to “Installing the RDAC failover driver” on page 105 or “Installing the DMP failover driver” on page 107. You do not need to modify JNI HBA settings if you have a direct-attached configuration.

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

Before you begin: Table 23 on page 96 lists the configuration file for each JNI adapter driver package. Also, some of the steps in these procedures refer to Appendix B, “JNI and QLogic host bus adapter settings,” on page 149.

Complete the following procedures to modify the JNI HBA settings.

Modifying the loop settings:

1. Run the following commands.

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

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

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

2. In the Vi Editor, uncomment and modify the loop attributes using the information in Appendix B, “JNI and QLogic host bus adapter settings,” on page 149.
3. Run the following command to save changes made to the JNI configuration file.

```
#:wq
```

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

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

Gathering the binding information:

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

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

The messages log displays.

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

```
# /SunOS
```

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

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

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

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

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

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

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

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

target0_hba variables:

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

Example:

```
targetZ_hba = "jnic146x0"
```

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

- b. Set the target to the second value that you recorded from the */var/adm/messages* file.

Example:

```
targetY_hba = "jnic146x0"
```

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

target_wwpn variables:

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

Example:

```
target_Z_wwpn="controller_1_wwpn"
```

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

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

Example:

```
target_Y_wwpn="controller_2_wwpn"
```

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

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

```
#:wq
```

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

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

Modifying the JNI HBA settings for a Brocade 2 Gb switch

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

If you are using any other switch, see “Modifying the JNI HBA settings for McData, Cisco, InRange, and Brocade 1 GB switches” on page 97 for the correct modification procedures.

Complete the following procedures to modify the JNI HBA settings.

Modifying the loop settings:

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

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

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

```
# /Sun0S
```

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

```

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

```

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

```

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

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

```

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

```

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

```

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

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

Before you begin: See Appendix B, “JNI and QLogic host bus adapter settings,” on page 149 and modify settings accordingly.

1. Type the following command:

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

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

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


3. Paste the copy immediately following the text that you copied in step 2 on page 100, as shown in the following example:

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

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

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

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

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

Note: Brocade 2 Gb switches (public):

- FcLoopEnable=1
- FcFabricEnable=1

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

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

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

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

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

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

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

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

#type: string
#Default

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

#####
```

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

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

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

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

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

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

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

- Type the following command to save the changes to the JNI configuration file:

```
#:wq
```

- Type the following command to reboot and configure the Solaris host:

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

When you are finished, install the failover driver on the host. See “Installing the RDAC failover driver” on page 105 or “Installing the DMP failover driver” on page 107 for more information.

Installing the QLogic adapter driver package

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

Note: The QLogic HBA is supported only on Solaris 8 and 9 operating systems.

Before you begin:

- Ensure that the QLogic HBAs have been installed, and are attached either directly to the controllers or to the SAN fabric switch.
- If the HBAs are attached to a SAN fabric switch, ensure that the zones for the DS4000 storage subsystem are created and enabled.

Notes:

1. All fabric switches must be zoned in such a way that a single HBA can access only one controller per storage array.
2. For more information about zoning and enabling zones, see Appendix E, “Connecting HBAs in a FC switch environment,” on page 165.

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

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

www.QLogic.com

2. Un-compress the QLogic adapter driver package, by typing the following command:

```
#uncompress qlogic_adapter_driver.
```

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

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

```
#pkgadd -d adapter_driver_pkg
```

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

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

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

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

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

Note: When the following screen displays, press **Enter** to install the driver object in the default directory (*/kernel/drv*):

```
Where do you want the driver object installed?
[default = /kernel/drv]
```

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

```
#pkginfo -l pkg_version
```

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

7. The following screen displays. Type `y`, then press **Enter**.

```
Do you want to continue with the installation of the pkg_version?  
[y, n, ?]
```

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

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

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

When you are done, go to the next procedure, “Verifying the HBA installation.”

Verifying the HBA installation

Complete the following steps to verify that the HBAs are installed:

1. Type the following command:

```
# view /var adm messages
```

A list of administrative messages displays.

2. Scroll down the list until you see the HBAs listed.
3. Record the WWPN and the HBA number for each of the two HBA cards.

When you are done, go to the next procedure, “Modifying the QLogic HBA settings.”

Modifying the QLogic HBA settings

Complete the following steps to modify the QLogic HBA settings:

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

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

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

2. In the Vi Editor, uncomment and modify the loop attributes on each of the two QLogic HBA cards, using the information that is listed in Appendix B, “JNI and QLogic host bus adapter settings,” on page 149.
3. Save the changes to the QLogic configuration file, by typing the following command:

```
#:wq
```

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

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

When you are finished, install the failover driver, as described in the next section.

Installing a failover driver on Solaris

You can install either the RDAC failover driver or the DMP failover driver on your Solaris system.

This section includes the following procedures:

- “Installing the RDAC failover driver”
- “Installing the DMP failover driver” on page 107

Installing the RDAC failover driver

This section describes how to install RDAC on a Solaris host.

Before you begin: You must install a JNI or QLogic driver package before you install RDAC. If you have a SAN-attached configuration, you must also modify the JNI or QLogic configuration file before you install RDAC. If you fail to follow the procedures in this order, problems can occur.

For information about how to install the JNI or QLogic HBA driver packages and modify the settings, see “Configuring the host bus adapter cards” on page 96.

Note: Modifying failover settings in the JNI or QLogic configuration file after installing RDAC requires the removal of the RDAC from the host.

Steps for installing the RDAC failover driver

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

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for Solaris and click the link in the Current Version column.

The latest versions are listed in the readme file.

Complete the following steps to install RDAC:

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

Note: In this procedure, the installation CD is mounted at `/cdrom/SM91`. Modify these commands as needed for your installation.

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

```
# pkgadd -d path/filename.pkg
```

where *path/filename* is the directory path and name of the package that you want to install.

The installation process begins.

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

```
The following packages are available:

1 RDAC                               Redundant Disk Array Controller
                                      (sparc) version number

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

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

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

- Type **y** and press **Enter**.
- 5. The following prompt is displayed:

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

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

[y, n, ?]
```

- Type **y** and press **Enter**.
The installation process continues.
- 6. When the RDAC package has been successfully installed, the following message is displayed:

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

- Ensure that the variables in the configuration files for the JNI adapter cards have been set to the correct values. For information about these configuration files and their variables, see “Modifying the JNI HBA settings” on page 97.
- 7. Reboot the Solaris host by typing the following command:

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

Attention: Any modifications to the persistent bindings in the `jnic146.conf` requires the removal of RDAC. After the RDAC is removed you can modify the persistent bindings in the `jnic146x.conf`, and then re-install RDAC.

Complete the following steps to modify the `sd.conf` or `jnic146x.conf` files:

- 1. Remove RDAC by typing the following command:

```
#pkgrm RDAC_driver_pkg_name
```

where `RDAC_driver_pkg_name` is the name of the RDAC driver package that you want to remove.

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

```
#pkginfo RDAC_driver_pkg_name
```

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

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

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

4. Modify persistent bindings in the *jnic146x.conf* or edit the *sd.conf* file by typing the following command:

```
# vi /kernel/drv/jnic146x.conf or sd.conf
```

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

```
#:wq
```

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

```
#pkgadd -d RDAC_driver_pkg_name
```

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

6. Verify package installation by typing the following command:

```
#pkginfo RDAC_driver_pkg_name
```

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

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

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

Note: You must reboot the host after modifying the *jnic146x.conf* file, because the *jnic146x.conf* driver is only read during the boot process. Failure to reboot the host might result in some devices being inaccessible.

Installing the DMP failover driver

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

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

www.veritas.com

System requirements

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

- Solaris operating system
- VERITAS Volume Manager 3.5
- *SMibmasl.pkg* (enables Solaris to recognize the DS4000 machine type)

DMP installation overview

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

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

Note: For more information about zoning and enabling zones, see Appendix E, “Connecting HBAs in a FC switch environment,” on page 165.

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

1. “Preparing for VERITAS DMP installation”
2. “Installing VERITAS Volume Manager packages for DMP” on page 110
3. “Installing the SMibmasl software package” on page 113
4. “Configuring VERITAS Volume Manager” on page 113
5. “Starting the VERITAS Enterprise Administrator” on page 117

Preparing for VERITAS DMP installation

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

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

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

Notes:

- Each target represents a controller to a subsystem, and each LUN represents a logical drive.
- If you are adding additional target or LUN definitions to the `/kernel/drv/sd.conf` file for an existing DMP configuration, be sure to reboot the Solaris host.
 - a. Open the `/kernel/drv/sd.conf` file with the Vi Editor, by typing the following command:

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

The file looks similar to the following example:


```

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

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

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

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

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

```

- b. Add additional target and LUN definitions, using the Vi Editor. In the following example, it is assumed that the Solaris host is attached to one DS4000 subsystem with three LUNs mapped to the DS4000 storage partition. In addition, the access LUN must be mapped to the partition.

```

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

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

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

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

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

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

```

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

```
# :wq
```

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

```
# pkginfo -l RDAC
```

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

```
# pkgrm RDAC
```

5. Verify that a host partition has been created. If not, create one by following the procedure that is described in “Storage Partitioning: Defining a Solaris host group” on page 93.

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

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

```
# SMdevices
```

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

```
# format
```

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

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

Installing VERITAS Volume Manager packages for DMP

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

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

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

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

```
#df
```

The result is similar to the following example:

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

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

```
# cd /cdrom/storage_solutions_solaris_3.5dc1
```

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

```
# ls
```

```
# ./installer
```

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

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

4. A new screen displays. Type **y** and press **Enter** to install VRTSvlic, as shown in the following example:

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

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

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

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

1= Add License Key 2=Installation Menu 3=Refresh h=help
p= Product Description q=quit
Enter [1,2,3,h,p,q]: 1
```

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

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

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

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

VERITAS Volume Manager 3.5 installation documentation if you want to install the VERITAS Volume Manager GUI.

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

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

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

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

VERITAS Storage Solutions
Available products:
<ul style="list-style-type: none"> 1) Volume Manager 2) File System 3) Foundation Suite (Quick Start, HA & Flash Snap) 4) Volume Replicator 5) Database Edition for Oracle 6) Cluster Server 7) SanPoint Foundation Suite 8) Cluster Server QuickStart 9) Database Edition/Advance Cluster for Oracle 9i 10) Cluster Server Traffic Director
<ul style="list-style-type: none"> q) Return to the main menu h) Installation help
Enter the number of the product to install or [q,h]: 1
VERITAS Storage Solutions
You are about to begin installation of Volume Manager
Do you want to continue? [y,n]: y

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

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

10. Exit the VERITAS installer by selecting option q (Return to the main menu) on the products screen.
11. A product summary screen displays. Exit the screen by selecting option q (quit), as shown in the following example:

VERITAS Storage Solutions		
VERITAS Product	Currently Installed	Licensed
=====		
Volume Manager	3.5, REV=06.21.2002.23.14	Yes
File System	3.5, REV=GA06	Yes
Foundation Suite	3.5	Yes
Volume Replicator	3.5, REV=06.21.2002.23.14	Yes
Flash Snap	3.5, REV=06.21.2002.23.14	Yes
Database Edition for Ora	Not Installed	No
1= Add License Key 2=Installation Menu 3=Refresh h=help p= Product Description q=quit		
Enter [1,2,3,h,p,q]: q		

Installing the SMibmasl software package

Complete the following steps to install the SMibmasl software package:

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

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

```
#pkgadd -d SMibmasl_pkg
```

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

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

Configuring VERITAS Volume Manager

Complete the following steps to configure the VERITAS Volume Manager:

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

```
# vxinstall
```

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

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

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

2. The Volume Manager Installation screen displays. Choose whether to set up enclosure-based names for the DS4000-attached disks by typing y (yes) or n (no), as shown in the following example:

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

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

enclosurename_diskno

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

Some examples would be:

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

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

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

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

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

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

Enclosure: Disk IBM_FAStT0

Hit RETURN to continue
```

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

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

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

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

Select an operation to perform: 2
```

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

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

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

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

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

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

Disk array serial number: DISKS

Generating list of disks in disk array Disk ...

<excluding root disk c0t0d0>

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

c0t1d0

Hit RETURN to continue.
```

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

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

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

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

Select an operation to perform: 3
```

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

```

Volume Manager Custom Installation
Menu: VolumeManager/Install/Custom/disk/IBM_FASt0

Disk array serial number : 600A0B800000F14400000

Generating list of disk in disk array IBM_FASt0 ....

    The Volume Manager has detected the following
disk in enclosure IBM_FASt0:

    C8t4d0  c8t1d1  c8t4d2  c8t4d3

Hit RETURN to continue.

```

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

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

```

Installation options for disk c8t4d0
Menu: VolumeManager /Install/Custom/Disk/IBM_FASt0

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

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

Select an operation to perform: 2

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

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

```

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

```

Installation options for disk c8t4d1
Menu: VolumeManager /Install/Custom/Disk/IBM_FASt0

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

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

Select an operation to perform: 3

```

13. Be sure to repeat step 12 for each of the remaining disks on the DS4000 storage array.

14. A summary screen displays, similar to the following example. Press **Enter**.

```
Volume Manager Custom Installation
Menu: VolumeManager /Install/Custom/Disk/IBM_FASTT0

The following is a summary of your choices.

C8t4d0    New disk

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

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

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

Starting the VERITAS Enterprise Administrator

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

Tip: VEA starts automatically after a reboot.

Complete the following steps to start VEA:

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

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

2. Restart the VEA service, by typing the following command:

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

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

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

4. Start the VEA GUI, by typing the following command:

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

5. Login to the VEA GUI, and re-scan.

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

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

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

Verifying external storage

After you install the failover driver, complete the following steps to verify that the host recognizes the external storage.

Before you begin: SMutil must be installed on the host to enable the **SMdevices** and **hot_add** commands, which you use in the following procedure. For information about how to install SMutil, see “Steps for installing SMutil software (required)” on page 90.

Restriction: Do not delete the access LUN or the access logical drive. RDAC uses the access LUN for communication between controllers and DS4000 Storage Manager. Therefore, a storage subsystem with two controllers would require two access LUNs. The access LUN is typically assigned the LUN number 31, by default.

1. Run the following command to verify that the host recognizes the external DS4000 storage:

```
# /opt/IBM_FAStT/util/SMdevices
```

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

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

2. Run the following command to verify that SMutil and the host both recognize all the configured storage:

```
# format
```

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

- Recheck the zones.
- Ensure that the zones have been enabled.
- Recheck the targets and WWPNs.
- Check to see if the DS4000 Controllers and storage expansion enclosures are powered on.
- Check to see if the switch or switches are powered on.
- Check Fibre Channel cables.

Redistributing logical drives in case of failure

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

Complete the following steps to manually redistribute logical drives to their preferred paths:

1. Repair or replace any faulty components. For more information, see the *IBM TotalStorage DS4000 Installation and Support Guide* for the appropriate DS4000 storage server.
2. Using the Subsystem Management window, redistribute logical drives to their preferred paths by clicking **Advanced** → **Recovery** → **Redistribute Logical Drives**.

Chapter 5. Using DS4000 with Linux on POWER-based hosts

This chapter provides the following information:

- “Overview”
- “DS4000 hardware and firmware requirements” on page 122
- “Restrictions” on page 124
- “Configuring storage subsystems on POWER-based Linux hosts” on page 125
- “Creating direct and SAN-attached configurations” on page 132
- “Configuring the host bus adapter drivers” on page 133
- “Performing initial device identification” on page 134

For host operating system requirements, see “SUSE Enterprise Linux Server system requirements” on page 124

Overview

The DS4000 Storage Manager client software is not currently available for POWER-based SUSE Linux hosts.

To manage DS4000 storage subsystems with your POWER-based SUSE Linux host, you must install the Storage Manager client software (SMclient) on an AIX 5L, Solaris, HP-UX, Windows or i386 processor-based Linux server, and use that server as your storage management station.

Note: For information on installing DS4000 Storage Manager on an AIX, Solaris, or HP-UX management station, see the appropriate installation chapter in this document:

- Chapter 2, “Installing and configuring storage management software on AIX host systems,” on page 23
- Chapter 3, “Installing and configuring storage management software on HP-UX host systems,” on page 63
- Chapter 4, “Installing and configuring storage management software on Solaris host systems,” on page 83

For information on installing DS4000 Storage Manager on a Windows or i386 processor-based Linux management station, see *IBM TotalStorage DS4000 Storage Manager 9 Installation and Support Guide for Windows 2000/Server 2003, NetWare, ESX Server, and Linux (GC26-7706)*.

For Linux on POWER/pSeries support, see the following Web sites:

Linux on POWER resource center

www.ibm.com/servers/enable/linux/power/

Linux on pSeries support

www.ibm.com/servers/eserver/pseries/linux/

Hardware, software and host system requirements

This section lists the the following requirements:

- Hardware and firmware requirements for the DS4000 configuration
- Hardware and software requirements for the SUSE Linux host system

Note: For the latest installation requirements and user information about Storage Manager and hardware components, see the readme file that is located in the Linux on POWER DS4000 Installation CD. The most recent copy of the readme file, which supersedes this document, is maintained at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for Linux on POWER and click the link in the Current Version column.

DS4000 hardware and firmware requirements

Table 24 lists the DS4000 Fibre Channel storage servers that you can use with POWER-based Linux hosts.

Note: To ensure that you have the latest versions of the firmware and the NVSRAM, see the latest readme file at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the resulting download files matrix, the links to the firmware downloads also provide the readme files for Storage Manager.

Table 24. Hardware that is supported with POWER-based SUSE Linux hosts

Hardware	Product Name	Model	Minimum firmware version
storage expansion enclosures storage expansion enclosure	IBM TotalStorage DS4000 EXP100 SATA storage expansion enclosure	1710-10U or 10X	9554
	IBM TotalStorage DS4000 EXP710 Fibre Channel storage expansion enclosure		
	IBM TotalStorage DS4000 EXP700 Fibre Channel storage expansion enclosure	1740-1RU or 1RX	9326
	IBM TotalStorage DS4000 EXP500 Fibre Channel storage expansion enclosure	3560-1RU or 1RX	9166

Table 24. Hardware that is supported with POWER-based SUSE Linux hosts (continued)

Hardware	Product Name	Model	Minimum firmware version
Host Bus Adapters	IBM HA Emulex LP9000	FC 6228	3.91a1
	IBM HA Emulex LP9002-F2	FC 6239	1.81x1
	IBM HA Emulex LP10000	FC 5716	1.90x4
Switches	Brocade switch	2109-S08 2109-S16	2.6.2a
	Brocade switch	2109-F16	3.1.2a
	Brocade switch	2109-M12	4.0.2c
	Brocade switch	2109-F32	4.2.0c
	Brocade switch	2109-F08	3.1.2a
	McData switch	2032-064, 2032-140	6.01.00 - 18
	McData switch	2031-016	6.01.00 - 18
	McData switch	2031-216, 2031-224	6.01.00 - 18
	InRange switch	2042-001	4.3.0
Cisco switch	2062-D01, 2062-D07, 2062-T07	1.3.4a	

Table 25 lists the DS4000 EXP storage expansion enclosures and other hardware that are supported with POWER-based Linux hosts.

Note: For important information about using the DS4000 EXP100 SATA storage expansion enclosure, see *IBM TotalStorage DS4000 EXP100 storage expansion enclosure Installation and User's Guide* (GC26-7601).

Table 25. Hardware that is supported with POWER-based SUSE Linux hosts

Hardware	Product Name	Model	Minimum firmware version
storage expansion enclosuresDS4000	IBM TotalStorage DS4000 EXP100SATA storage expansion enclosure	1710-10U or 10X	9554
	IBM TotalStorage DS4000 EXP710Fibre Channel storage expansion enclosure		
	IBM TotalStorage DS4000 EXP700 Fibre Channel storage expansion enclosure	1740-1RU or 1RX	9326
	IBM TotalStorage DS4000 EXP500 Fibre Channel storage expansion enclosure	3560-1RU or 1RX	9163

Table 25. Hardware that is supported with POWER-based SUSE Linux hosts (continued)

Hardware	Product Name	Model	Minimum firmware version
Host Bus Adapters	IBM HA Emulex LP9000	FC 6228	3.91
	IBM HA Emulex LP9002-F2	FC 6239	1.81x1
	IBM HA Emulex LP10000	FC 5716	1.81a1
Switches	Brocade switch	2109-F16	3.12a

Recommendation: When you set up the hardware, connect the HBAs last. Do not connect them before the storage servers are set up and the system is loaded.

SUSE Enterprise Linux Server system requirements

This section lists the minimum hardware and software requirements for a POWER-based host running SUSE Linux in a DS4300 Turbo, DS4400, or DS4500 storage subsystem that is managed by DS4000 Storage Manager.

Hardware requirements

The following hardware is required at minimum:

- **Host:** IBM @server iSeries™ or pSeries with 64-bit PowerPC Architecture™
- **Management station:**
 - For Windows or i386-based Linux management stations, see the minimum requirements listed for those operating systems in *IBM TotalStorage DS4000 Storage Manager 9 Installation and Support Guide for Windows 2000/Server 2003, NetWare, ESX Server, and Linux (GC26-7706)*.
 - For AIX, HP-UX or Solaris management stations, see the minimum requirements listed in the appropriate chapters in this document.

Software requirements

The following software distributions must be the specified versions:

SUSE LINUX Enterprise Server 9

- **Version:** RC5
- **Kernel:** 2.6.5 – 7.97

Restrictions

The following restrictions apply to POWER-based Linux host configurations:

SAN and connectivity restrictions

- POWER-based SUSE LINUX Enterprise Server 9 hosts can support one host bus adapter (HBA) per controller. You can configure one HBA to each mini-hub on the DS4000 controller. Each HBA-to-controller relationship must reside in a separate zone.
- Direct-attach configurations are restricted to single-initiator configurations only. You can only make one connection to each minihub on a DS4000 storage server.

Note: DS4300 Turbo storage servers do not have minihubs.

- Single-switch configurations are allowed, but each HBA and DS4000 controller combination must be in a separate SAN zone.

- Other storage devices, such as tape devices or other disk storage, must be connected through separate HBAs and SAN zones.

Partitioning restrictions

- All logical drives that are configured for Linux on POWER must be mapped to a Linux host group. For more information, see “Storage Partitioning premium feature” on page 21.
- On each controller, you must configure at least one LUN with an ID between 0 and 31 that is not a UTM or access logical drive.

Interoperability restrictions

- You cannot use F-RAID Manager.

Configuring storage subsystems on POWER-based Linux hosts

If you have purchased a Storage Partitioning premium feature, ensure that the premium feature is enabled. For more information, see “Enabling your premium features” on page 21.

Note: DS4000 Storage Manager software is not currently available for SUSE Linux on POWER-based hosts. Complete the DS4000 Storage Manager procedures in this section using SMclient on your storage management station, as described in the “Overview” on page 121.

Adding storage subsystems to DS4000 SMClient

Perform the following steps to add storage subsystems to DS4000 SMclient:

1. The subsystem must be physically configured for direct management using Ethernet connections from the management station to each controller. Install SMclient on the management station before configuring the subsystem.

Note: See “Setting up IP addresses for DS4000 storage controllers” on page 17 for information about assigning initial IP addresses to the controllers.

2. After the disk subsystem is configured on the network, start SMclient on the management station, as follows:

AIX management station:

Type the following command:

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

HP-UX or Solaris management station:

Type the following command:

```
# SMclient
```

Windows management station:

Select **Start** → **Programs** → **IBM FASTT Storage Manager 9 Client**.

The Enterprise Management window opens.

3. Complete the following steps to add the storage subsystems to the SMclient:
 - a. In the Enterprise Management window, click **Edit** → **Add Device**.
 - b. In the Add Device window, type the IP address of the first controller in the storage subsystem and click **Add**.

- c. Type the IP address of the second controller and click **Add**, and then click **Close**.

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

Upgrading DS4000 firmware and NVSRAM

The procedures in this section describe how to download and install DS4000 controller firmware and NVSRAM.

Before you begin: Read the information in the following sections before you begin the download procedures:

- “Checking current versions of firmware and NVSRAM”
- “Staged download” on page 127

Checking current versions of firmware and NVSRAM

Versions of firmware and NVSRAM that are provided with this product might be later releases than those described in this document. Check your current versions and the latest readme file to find out if you have the latest versions.

To find whether you have the latest versions of firmware and NVSRAM, complete the following steps:

1. In the Subsystem Management window, click **Storage Subsystem** → **View Profile**. The Storage Subsystem Profile window displays.
2. In the Storage Subsystem Profile window, click the **Controllers** tab to see the current versions of firmware and NVSRAM.
3. Click **Close** to close the Storage Subsystem Profile window.
4. To find whether the versions of firmware and NVSRAM that are currently on your storage subsystem are the latest versions, see the readme file at the following Web site:

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

When the page opens, click the **Download** tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for Linux on POWER and click the link in the **Current Version** column.

The latest versions are listed in the readme file.

5. If you find that you do not have the latest versions of the firmware and NVSRAM, click the **Back** button on your browser to return to the previous page. You can download the correct versions to a directory by clicking the appropriate link that is listed in the IBM DS4000 Controller Firmware matrix.

Traditional or staged firmware download

There are two methods for downloading the controller firmware:

- “Traditional download”
- “Staged download” on page 127

You can select which method you want to use during the upgrade process, as described in “Upgrading DS4000 controller firmware” on page 127 and “Upgrading the NVSRAM” on page 128.

Traditional download: The traditional download process takes significantly longer and must be done in one phase, rather than in two phases as with the staged

controller firmware download. Therefore the preferred method of controller firmware download is the staged controller firmware download.

Staged download: The staged controller firmware download feature separates downloading and activation into two separate steps.

Notes:

1. You cannot upgrade Storage Manager 8.4 (5.4x.xx.xx) firmware to Storage Manager 9.1 (6.1x.xx.xx) firmware using staged download. You can only use staged download if the controller firmware version is already 6.1x.xx.xx, and you are downloading another version of 6.1x.xx.xx firmware.
2. Staged controller firmware download is not supported on DS4400 storage servers.

You can download controller firmware from the storage management software to all storage subsystem controllers for later activation.

The firmware download process occurs in two distinct phases: firmware transfer and firmware activation. You can perform the time-consuming task of loading the firmware online so that it is functionally transparent to the application. You can then defer the activation of the loaded firmware to a convenient time.

Some reasons for activating firmware files at a later time include the following:

- **Time of day** - Activating can take a long time, so you might want to wait until I/O loads are lighter because the controllers will go offline briefly to load the new firmware.
- **Type of package** - You might want to test the new firmware package on one storage subsystem before loading the files onto other storage subsystems.

Note: The ability to download firmware and activate it later depends on the version of storage management software and type of controller in the storage subsystem.

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

Upgrading DS4000 controller firmware

Complete the following steps to upgrade the controller firmware:

Before you begin: You must have management connections to both controllers, and both controllers must be in an optimal state before you start the controller firmware download.

1. Stop I/O to the controllers.
2. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.
3. Download the correct version of the firmware from the Web site that is listed in “Checking current versions of firmware and NVSRAM” on page 126. Place the file in a designated directory on the host system.
4. In the Subsystem Management window, click **Advanced** → **Maintenance** → **Download** → **Controller Firmware**. The Download Firmware window displays.

5. Click **Browse** to select the appropriate firmware file.
6. If you want to perform a staged firmware download, select the check box next to "Transfer files but don't activate them (activate later)." If you want to perform a traditional firmware download, do not select the check box, and the firmware will load immediately.
7. Click **OK** to update both controllers.
8. Click **Yes** to start the download. A new window opens.
9. Close the Subsystem Management window and then reopen it to complete the firmware update.
10. Verify that the firmware was successfully installed.
11. Restart I/O.

Upgrading the NVSRAM

Complete the following steps to upgrade NVSRAM. NVSRAM will be activated immediately.

Before you begin: You must have management connections to both controllers, and both controllers must be in an optimal state before you start the NVSRAM download.

1. Stop I/O to the controllers.
2. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.
3. Download the correct version of NVSRAM from the Web site listed in "Checking current versions of firmware and NVSRAM" on page 126. Place the file in a designated directory on the host system.
4. In the Subsystem Management window, click **Advanced** → **Maintenance** → **Download** → **Controller NVSRAM**. The Download NVSRAM window displays.
5. Click **Browse** to select the appropriate NVSRAM file.
6. Click **OK**. The Confirm Download window displays.
7. Click **Yes** to start the download.
8. Verify that the NVSRAM was successfully installed.
9. Restart I/O.

Storage Partitioning: Defining a host group

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

For more information about storage partitioning, see "Storage Partitioning premium feature" on page 21.

This section contains the following procedures:

1. "Steps for defining a host group" on page 129
2. "Steps for verifying that the host ports match the host" on page 129

Before you begin: Read the following information.

- Create the Linux host group at the storage subsystem level. Do not create host groups at the default group level.

Exception: If you are running a DS4300 configuration without partitioning enabled, you can use the default host group.

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

Steps for defining a 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**.

Note: Make sure that the storage subsystem is highlighted in the left panel of the Subsystem Management window. Do not click on Undefined Mappings.

3. Type the name of the new host group (for example, Linux SLES 9). Click **Add**, and then click **Close**.
4. Highlight the new host group and click **Mappings** → **Define** → **Host**.
5. Type the name of the host to which the storage subsystem is attached (for example, linuxserver). Click **Add**, and then click **Close**.
6. Highlight the host that you just added, then right-click and select **Define Host Port**.
7. Select the host port identifier for the first HBA (for example, 10:00:00:00:c9:24:0c:3f). If you do not see the identifier that you are looking for, see the note at the end of this procedure.
8. Change the host type to **Linux** (do not select Linux Cluster). Click **Add**.

Important: Failure to change the host type from the default to Linux might cause undesired results.

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

Note: The host port identifier is the WWPN of the host bus adapter. If you do not see the one you want in the host port identifier drop-down menu, check the following:

Fibre-channel cable

You might not be using the correct cable. Verify that you are using the correct fibre-channel cable, and that it is connected properly.

Switch zoning

If you are connected to a switch, ensure that it is zoned correctly.

lpfc module

You might not have initialized the lpfc module. Type the following command for the server that does not display in the menu:

```
# modprobe lpfcdd
```

If you still do not see the correct WWPNs, try one of the methods listed in “Viewing information about HBAs” on page 130.

Steps for verifying that the host ports match the host

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

1. Type the following command to view the HBAs:

```
lsdev | grep lpfcdd
```

A list that contains all the HBAs that are in the system is displayed.

2. In the list, identify the host number of the HBA that is connected to the DS4000.
3. Type the following command to view the WWPNs of the HBA and connected DS4000 controllers:

```
cat /proc/scsi/lpfc/host#
```

where *host#* is the host number of the HBA.

4. Verify that the network address number matches the host port number that displays in host partition table of the DS4000 SMclient.
5. Repeat this procedure to verify the other host ports.

Viewing information about HBAs

You can use one of the following methods to view information (such as WWPNs, firmware levels, and other properties) about the HBAs.

Viewing HBA information, Method 1:

1. Type the following commands:

```
# cd /proc/scsi/lpfc
# ls
```

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

```
. . . 10 11 7 8 9
```

2. Type the following command:

```
# cat hba#
```

where *hba#* is the number of an HBA in the list. The HBA's properties display, as shown in the following example:

```
e1m17c18:/proc/scsi/lpfc # cat 11
Emulex LightPulse FC SCSI 2.10f
HBA: Emulex LightPulse LP9002 (2 Gigabit) on PCI bus 2a device 08 irq 58
SerialNum: 1C21908CE5
Firmware Version: 3.82A1 (C2D3.82A1)
Hdw: 2002606d
VendorId: 0xf90010df
Portname: 10:00:00:00:c9:2d:29:f1   Nodename: 20:00:00:00:c9:2d:29:f1

Link Up - Ready:
  PortID 0x10200
  Fabric
  Current speed 2G

lpfc4t00 DID 010400 WWPN 20:02:00:a0:b8:13:0f:d3 WWNN 20:02:00:a0:b8:13:0f:d2
lpfc4t01 DID 010500 WWPN 20:0a:00:a0:b8:0f:33:79 WWNN 20:0a:00:a0:b8:0f:33:78
lpfc4t02 DID 010600 WWPN 20:08:00:a0:b8:12:aa:53 WWNN 20:08:00:a0:b8:12:aa:52
lpfc4t03 DID 010700 WWPN 20:04:00:a0:b8:12:db:f1 WWNN 20:04:00:a0:b8:12:db:f0
lpfc4t04 DID 010a00 WWPN 20:10:00:a0:b8:06:64:1e WWNN 20:10:00:a0:b8:06:64:1d
e1m17c18:/proc/scsi/lpfc #
```

Viewing HBA information, Method 2:

1. Type the following commands:

```
# cd /sys/class/scsi_host
# ls
```

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

```
. .. host0 host1 host10 host11 host7 host8 host9
```

2. Type the following commands:

```
# cd host#
# ls
```

where *host#* is the number of a host in the list. The next list of files displays, as shown in the following example:

```
. cmd_per_lun host_busy scan unchecked_isa_dma
.. device proc_name sg_tablesize unique_id
```

3. Type the following commands:

```
# cd device
# ls
```

A list of HBAs displays, as shown in the following example.

Note: [11:0:0:5] = [HBA : scsi_bus : scsi_id : LUN]

```
.          11:0:0:5  11:0:1:2  11:0:2:10  11:0:3:1  11:0:4:0  11:0:4:6
..         11:0:0:6  11:0:1:3  11:0:2:2  11:0:3:10 11:0:4:1  11:0:4:7
11:0:0:0   11:0:0:7  11:0:1:4  11:0:2:3  11:0:3:2  11:0:4:10 11:0:4:8
11:0:0:1   11:0:0:8  11:0:1:5  11:0:2:4  11:0:3:3  11:0:4:11 11:0:4:9
11:0:0:10  11:0:0:9  11:0:1:6  11:0:2:5  11:0:3:4  11:0:4:12 detach_state
11:0:0:11  11:0:1:0  11:0:1:7  11:0:2:6  11:0:3:5  11:0:4:13
11:0:0:12  11:0:1:1  11:0:1:8  11:0:2:7  11:0:3:6  11:0:4:2
11:0:0:2   11:0:1:10 11:0:1:9  11:0:2:8  11:0:3:7  11:0:4:3
11:0:0:3   11:0:1:11 11:0:2:0  11:0:2:9  11:0:3:8  11:0:4:4
11:0:0:4   11:0:1:12 11:0:2:1  11:0:3:0  11:0:3:9  11:0:4:5
```

Mapping LUNs to a DS4000 storage partition

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

Before you begin: Read “DS4000 Storage Manager performance planning” on page 18 for information about planning the number of LUNs to map to each partition.

1. Select the Mappings view of the Subsystem Management window.
2. In the topology tree, right-click the host on which you want to map LUNs, and select **Define Additional Mappings**. The Define Additional Mapping window opens.
3. In the Define Additional Mapping window, select the following options, and then click **Add**:

- Host group or host
- Logical unit number (LUN)(0-255)
- Logical drive

Repeat these steps for each LUN.

Notes:

1. It is recommended that you do not map at the host group level. If you do, every host can see all the LUNs.
2. Map the LUNs using consecutive numbers. For example, map LUNs to numbers 0; 1; 2; 3; 4; 5; and so on, without skipping any numbers. Failure to do so might cause undesired results.
3. Be sure to map LUN 0. Failure to do so might cause undesired results.

Creating direct and SAN-attached configurations

DS4000 Storage Manager supports IBM DS4000 Fibre Channel Storage Servers in direct-attached POWER-based Linux host configurations or in a SAN environment through switches in POWER-based Linux host configurations.

Creating a direct-attached configuration

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

- One or two POWER-based Linux host servers can be connected to the DS4000 storage servers.
 - No external hubs are being used.
 - Two-server DS4400 or DS4500 configurations require four host-side minihubs, each with exactly one fibre-channel connection from each HBA to a minihub.
- Note:** DS4300 Turbo storage servers do not have minihubs.
- Single HBA configurations are allowed.

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

1. Connect the HBAs to each controller or minihub port of the DS4000 storage server.
2. Configure and verify the configuration.

Creating a SAN-attached configuration

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

- Multiple fibre-channel HBAs within the same server cannot “see” the same DS4000 controller port.
- The IBM fibre-channel HBAs are isolated from each other if they are connected to the same switch that is connected to the same DS4000 controller port.
- Each fibre-channel HBA and controller port must be in its own fabric zone, if they are connecting through a single fibre-channel switch, such as 2109-F16.

See the documentation that is provided by the switch manufacturer for more information about zoning. Multiple DS4000 devices can be configured to the same set of fibre-channel HBAs through a fibre-channel switch.

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

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

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

1. Connect the HBAs to the switch or switches.
2. Connect the DS4000 storage subsystems to the switch or switches.
3. Set the required zones on the fibre-channel switch or switches, if applicable.

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

4. Configure, then verify the configuration by logging into the switch to view the system.

Configuring the host bus adapter drivers

Each Emulex HBA that is used with POWER-based Linux hosts needs two separate HBA and utility driver packages. The standard packages are loaded onto the system with the SUSE LINUX Enterprise Server 9 download.

Use the **lsmod** command to find out whether the HBA device driver (lpfcdd module) is on the system.

If you do not intend to modify the driver, you can use the standard packages. Otherwise, download the latest packages and configure them using the procedures in this section.

Downloading the HBA driver and utility packages

Download the following packages from the installation CD, from the IBM DS4000 support Web site, or from the Emulex Web site:

	Compressed package filename	Uncompressed filename	Description
Utility package	EmulexLibApp_ <i>package_name</i> .tar.gz	emlslib (lputil)	Monitors HBAs and facilitates downloading of HBA firmware
Driver package	lpfcdriver- <i>package_name</i> .tar.gz	lpfc driver	Compiles changes in the configuration file (/etc/lpfc.conf)

Note: For the compressed file name, *package_name* refers to the firmware level of the package. For example: EmulexLibApp_2.10f.tar.gz

For each package, type the following command to uncompress the files:

```
# tar -z -xvf package_name.tar.gz
```

Configuring HBA drivers

Complete the following steps to set up the HBA drivers:

1. Create the SAN switch connections.
2. Type the following command to start the HBA driver:

```
# modprobe lpfcdd
```

The driver searches for all HBAs, and their connections, that are on the system.

3. Type the following command to view the **lpfcdd** output:

```
# lsscsi
```

A list of HBAs and their connections displays, as shown in the following example.

Note: [1:0:0:0] = [HBA : scsi_bus : scsi_id : LUN]

Column 4 shows the associated DS4000 type, and column 6 shows the Linux operating system device identifier.

```
e1m17c27:~ # lsscsi
[0:0:5:0] disk IBM ST336607LC C50H /dev/sda
[0:255:255:255]no dev IBM 570B001 0150 -
[1:0:0:0] disk IBM 1742-900 0520 /dev/sdb
[1:0:0:1] disk IBM 1742-900 0520 /dev/sdc
[1:0:0:2] disk IBM 1742-900 0520 /dev/sdd
[1:0:0:3] disk IBM 1742-900 0520 /dev/sde
[1:0:0:4] disk IBM 1742-900 0520 /dev/sdf
[1:0:0:5] disk IBM 1742-900 0520 /dev/sdg
[1:0:0:6] disk IBM 1742-900 0520 /dev/sdh
[1:0:0:7] disk IBM 1742-900 0520 /dev/sdi
[1:0:0:8] disk IBM 1742-900 0520 /dev/sdj
[1:0:0:9] disk IBM 1742-900 0520 /dev/sdk
[1:0:1:0] disk IBM 1742 0520 /dev/sdl
[1:0:1:1] disk IBM 1742 0520 /dev/sdm
[1:0:1:2] disk IBM 1742 0520 /dev/sdn
[1:0:1:3] disk IBM 1742 0520 /dev/sdo
[1:0:1:4] disk IBM 1742 0520 /dev/sdp
```

4. Type the following command to view the properties of each HBA (such as firmware version, controller WWPNs, and connection speed):

```
# cat /proc/scsi/lpfc/HBA#
```

Repeat for each HBA. (See an example of the **cat** command output in “Viewing information about HBAs” on page 130.)

When you are done, check the Mappings View of the Subsystem Management window to verify that the LUNs that are detected by the HBAs are correct.

Performing initial device identification

Identify the devices by using one of the following methods:

Device identification: Method 1

1. Type the following commands:

```
# rmmod lpfcdd
# modprobe lpfcdd
# lsscsi
```

A list of HBAs and their connections displays, as shown in the following example.

Note: [1:0:0:0] = [HBA : scsi_bus : scsi_id : LUN]

Column 4 shows the associated DS4000 type, and column 6 shows the Linux OS device number.

```
e1m17c27:~ # lsscsi
[0:0:5:0] disk IBM ST336607LC C50H /dev/sda
[0:255:255:255]no dev IBM 570B001 0150 -
[1:0:0:0] disk IBM 1742-900 0520 /dev/sdb
[1:0:0:1] disk IBM 1742-900 0520 /dev/sdc
[1:0:0:2] disk IBM 1742-900 0520 /dev/sdd
[1:0:0:3] disk IBM 1742-900 0520 /dev/sde
[1:0:0:4] disk IBM 1742-900 0520 /dev/sdf
[1:0:0:5] disk IBM 1742-900 0520 /dev/sdg
[1:0:0:6] disk IBM 1742-900 0520 /dev/sdh
[1:0:0:7] disk IBM 1742-900 0520 /dev/sdi
[1:0:0:8] disk IBM 1742-900 0520 /dev/sdj
[1:0:0:9] disk IBM 1742-900 0520 /dev/sdk
[1:0:1:0] disk IBM 1742 0520 /dev/sdl
[1:0:1:1] disk IBM 1742 0520 /dev/sdm
[1:0:1:2] disk IBM 1742 0520 /dev/sdn
[1:0:1:3] disk IBM 1742 0520 /dev/sdo
[1:0:1:4] disk IBM 1742 0520 /dev/sdp
```

2. Type the following commands:

```
# cd /proc/scsi/lpfc
# ls
```

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

```
. . . 10 11 7 8 9
```

3. For each HBA in the directory, type the following command:

```
# cat hba#
```

where *hba#* is an HBA in the directory. The HBA's properties display, as shown in the following example:

```
e1m17c18:/proc/scsi/lpfc # cat 11
Emulex LightPulse FC SCSI 2.10f
HBA: Emulex LightPulse LP9002 (2 Gigabit) on PCI bus 2a device 08 irq 58
SerialNum: 1C21908CE5
Firmware Version: 3.82A1 (C2D3.82A1)
Hdw: 2002606d
VendorId: 0xf90010df
Portname: 10:00:00:00:c9:2d:29:f1 Nodename: 20:00:00:00:c9:2d:29:f1

Link Up - Ready:
  PortID 0x10200
  Fabric
  Current speed 2G

lpfc4t00 DID 010400 WWPN 20:02:00:a0:b8:13:0f:d3 WWNN 20:02:00:a0:b8:13:0f:d2
lpfc4t01 DID 010500 WWPN 20:0a:00:a0:b8:0f:33:79 WWNN 20:0a:00:a0:b8:0f:33:78
lpfc4t02 DID 010600 WWPN 20:08:00:a0:b8:12:aa:53 WWNN 20:08:00:a0:b8:12:aa:52
lpfc4t03 DID 010700 WWPN 20:04:00:a0:b8:12:db:f1 WWNN 20:04:00:a0:b8:12:db:f0
lpfc4t04 DID 010a00 WWPN 20:10:00:a0:b8:06:64:1e WWNN 20:10:00:a0:b8:06:64:1d
e1m17c18:/proc/scsi/lpfc #
```

4. Check the information to verify connectivity with the storage controllers.

Device identification: Method 2

1. Type the following commands:

```
# cd /sys/class/scsi_host/  
# ls
```

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

```
. .. host0 host1 host10 host11 host7 host8 host9
```

2. Type the following commands:

```
# cd host#  
# ls
```

where *host#* is the number of a host in the list. The next list of files displays, as shown in the following example:

```
elml7c27:/sys/class/scsi_host/host1 # ls  
. cmd_per_lun host_busy scan unchecked_isa_dma  
.. device proc_name sg_tablesize unique_id  
  
elml7c27:/sys/class/scsi_host/host1 # cd device  
elml7c27:/sys/class/scsi_host/host1/device # ls  
. 1:0:0:8 1:0:1:4 1:0:2:12 1:0:3:1 1:0:4:0 1:0:4:6  
.. 1:0:0:9 1:0:1:5 1:0:2:2 1:0:3:10 1:0:4:1 1:0:4:7  
1:0:0:0 1:0:1:0 1:0:1:6 1:0:2:3 1:0:3:2 1:0:4:10 1:0:4:8  
1:0:0:1 1:0:1:1 1:0:1:7 1:0:2:4 1:0:3:3 1:0:4:11 1:0:4:9  
1:0:0:2 1:0:1:10 1:0:1:8 1:0:2:5 1:0:3:4 1:0:4:12 detach_state  
1:0:0:3 1:0:1:11 1:0:1:9 1:0:2:6 1:0:3:5 1:0:4:13  
1:0:0:4 1:0:1:12 1:0:2:0 1:0:2:7 1:0:3:6 1:0:4:2  
1:0:0:5 1:0:1:13 1:0:2:1 1:0:2:8 1:0:3:7 1:0:4:3  
1:0:0:6 1:0:1:2 1:0:2:10 1:0:2:9 1:0:3:8 1:0:4:4  
1:0:0:7 1:0:1:3 1:0:2:11 1:0:3:0 1:0:3:9 1:0:4:5
```

3. Compare the LUNs that are displayed in the **Isscsi** output to the LUNs that are displayed in the Mappings View of the Subsystem Management window.
4. If the number of LUNs displayed in the Subsystem Management window is different from the number of LUNs displayed in the **Isscsi** output, the most likely cause is that the LUNs are not mapped properly. In this case, ensure that LUN 0 is mapped to the host, and that the LUNs are mapped in consecutive order.

Device identification: Method 3

Complete the following steps to view HBA and DS4000 controllers' port connectivity using the switch interface:

1. Log in to the switch over the browser. (For example: <http://9.47.72.152>)
2. In the switch view, click **Admin**.
3. When prompted, log in as **root**. The Admin window opens.
4. In the Admin window, select **View**.
5. A report showing port status, zones and connections is displayed, as shown in Figure 13 on page 137 and Figure 14 on page 138.

Switch Information Report for IBM_2109_F16

List of Switches

Switch ID	Worldwide Name	Enet IP Addr	FC IP Addr	Name
1: fffc01	10:00:00:60:69:51:b5:48	9.47.72.152	0.0.0.0	>"IBM_2109_F16"

Ethernet IP Address: 9.47.72.152
Ethernet Subnet Mask: 255.255.255.0
Fibre Channel IP Address: none
Fibre Channel Subnet Mask: none
Gateway Address: 9.47.72.1

Kernel: 5.4
Fabric OS: v3.1.2a
Made on: Wed Mar 31 21:16:31 PST 2004
Flash: Wed Mar 31 21:17:22 PST 2004
BootProm: Tue Oct 30 10:24:38 PST 2001

List of Inter-Switch Links

Local Domain ID: 1

Local Port	Domain	Remote Port	State
------------	--------	-------------	-------

No ISL Ports !!

List of Ports

switchName: IBM_2109_F16
switchType: 9.2
switchState: Online
switchMode: Native
switchRole: Principal
switchDomain: 1
switchId: fffc01
switchWwn: 10:00:00:60:69:51:b5:48
switchBeacon: OFF
Zoning: ON (plinux)

port 0:	id N2 No_Light	
port 1:	id N2 No_Light	
port 2:	id N2 No_Light	
port 3:	id N2 No_Light	
port 4:	id N2 No_Light	
port 5:	id N2 No_Light	
port 6:	-- N2 No_Module	
port 7:	id N2 No_Light	
port 8:	-- N2 No_Module	
port 9:	id N2 Online	F-Port 10:00:00:00:c9:34:01:d2
port 10:	id N2 Online	F-Port 10:00:00:00:c9:37:73:f0
port 11:	id N2 Online	F-Port 20:04:00:a0:b8:12:6e:5f
port 12:	id N2 Online	F-Port 20:04:00:a0:b8:13:00:4a
port 13:	id N2 In_Sync	Disabled
port 14:	id N2 Online	F-Port 20:08:00:a0:b8:13:95:33
port 15:	id N2 No_Light	

Figure 13. Switch information report (part 1)

```

Name Server
{010900 010a00 010b00 010c00 010e00
5 Nx_Ports in the Fabric }

{
Type Pid      COS      PortName      NodeName      TTL(sec)
N 010900;    2,3;10:00:00:00:c9:34:01:d2;20:00:00:00:c9:34:01:d2; na
  FC4s: FCP
  Fabric Port Name: 20:09:00:60:69:51:b5:48
N 010a00;    2,3;10:00:00:00:c9:37:73:f0;20:00:00:00:c9:37:73:f0; na
  FC4s: FCP
  Fabric Port Name: 20:0a:00:60:69:51:b5:48
N 010b00;    3;20:04:00:a0:b8:12:6e:5f;20:04:00:a0:b8:12:6e:5e; na
  FC4s: FCP [IBM 1742 0520]
  Fabric Port Name: 20:0b:00:60:69:51:b5:48
N 010c00;    3;20:04:00:a0:b8:13:00:4a;20:04:00:a0:b8:13:00:48; na
  FC4s: FCP [IBM 1742-900 0520]
  Fabric Port Name: 20:0c:00:60:69:51:b5:48
N 010e00;    3;20:08:00:a0:b8:13:95:33;20:08:00:a0:b8:13:95:31; na
  FC4s: FCP [IBM 1722-600 0610]
  Fabric Port Name: 20:0e:00:60:69:51:b5:48

The Local Name Server has 5 entries }

Zoning Information
Defined configuration:
cfg: plinux red; blue
zone: blue 1,8; 1,9; 1,10; 1,11; 1,12; 1,13; 1,14; 1,15
zone: red 1,0; 1,1; 1,2; 1,3; 1,4; 1,5; 1,6; 1,7

Effective configuration:
cfg: plinux
zone: blue 1,8
1,9
1,10
1,11
1,12
1,13
1,14
1,15
zone: red 1,0
1,1
1,2
1,3
1,4
1,5
1,6
1,7

```

Figure 14. Switch information report (part 2)

Chapter 6. Completing the software configuration

After you finish installing DS4000 Storage Manager, you must complete your configuration. This chapter describes how to complete the DS4000 Storage Manager configuration by performing the following tasks:

- “Configuring storage subsystems”
- “Configuring alert notifications” on page 143
- “Renaming storage subsystems” on page 144
- “Creating a storage subsystem profile” on page 144
- “Performing other storage subsystem management tasks” on page 144

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

Overview of heterogeneous host configurations

The heterogeneous hosts feature enables hosts that are running different operating systems to access a single storage subsystem. DS4000 Storage Manager 9.1 supports up to 64 storage partitions, which enables a multiple host-type subsystem to share storage capacity, consolidate storage, and reduce storage management costs.

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

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

Before you begin setting up the configuration for you heterogeneous host, see the *IBM DS4000 Storage Manager Concepts Guide*.

Configuring storage subsystems

Use the following procedures to start the DS4000 SMclient from either the storage management station or from a host that is acting as a storage management station (a host with the client software installed). Use the SMclient to discover and configure each attached storage subsystem.

Performing an initial automatic discovery of storage subsystems

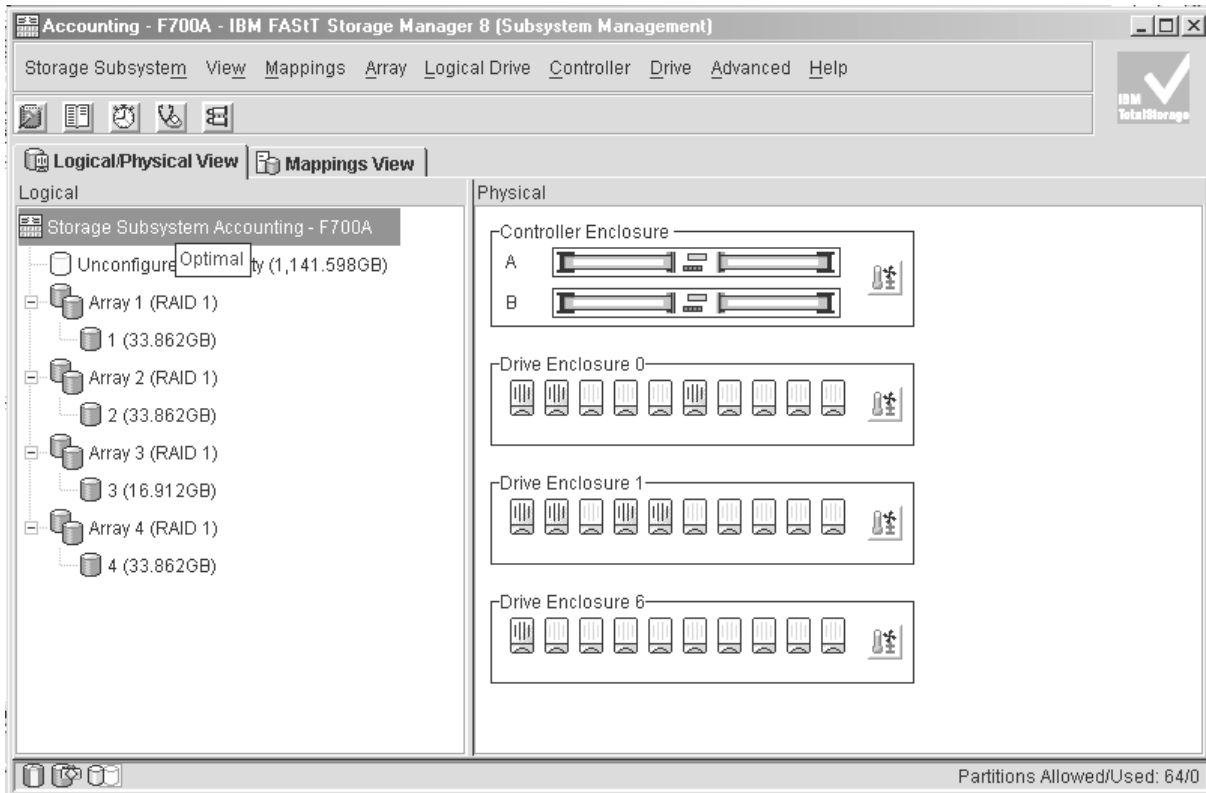
Complete the following steps to begin an initial automatic discovery of all the attached hosts and storage subsystems:

1. If you are running the DS4000 SMclient from AIX, first change to the /SMclient directory and then type SMclient. (For more information about initial automatic discovery on AIX systems, see “Identifying devices on AIX hosts” on page 38.) If you are running the client software from Solaris 8 or 9, type SMclient.

2. Click **Yes** to begin the initial automatic discovery.

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

When the initial automatic discovery is finished, all attached hosts and attached storage subsystems are displayed in the Enterprise Management window, as shown in Figure 15.



SJ001148

Figure 15. Enterprise Management window

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

Note: Failure to add both controllers results in a partially-managed device.
 - d. If a storage subsystem is duplicated in the device tree after an automatic discovery, remove the duplicate storage subsystem icon from the device tree by clicking **Edit** → **Remove Device** in the Enterprise Management window.
4. Ensure that the status of each storage subsystem is **Optimal**:

- If any device shows a status of **Unresponsive**, remove the device from the management domain, and then add it again. See the Enterprise Management window help topic on removing and adding devices.
 - If the device still shows an unresponsive status, contact your IBM technical-support representative.
5. Go to “Configuring alert notifications” on page 143.

Creating logical drives from free or unconfigured capacity

Perform following steps to create logical drives from free or unconfigured capacity nodes in the Subsystem Management window.

1. In the left panel of the Subsystem Management window, select either **Free Capacity** or **Unconfigured Capacity** and then click **Logical Drive —> Create**.

The Introduction (Create Logical Drive) window opens.

Note: If the Storage Partition feature is enabled, a Default Host Type window might open on top of the Introduction (Create Logical Drive) window, that asks you to select a default host type for the logical drive or drives that you are going to be create. Select the appropriate host type from the pull-down list and click **OK** to close the Default Host Type window.

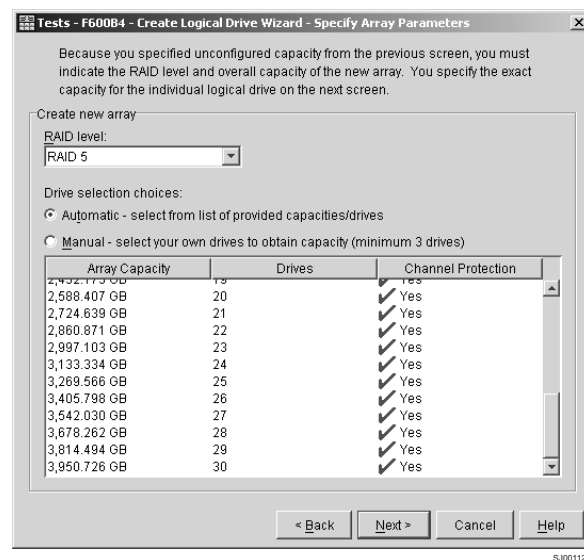


Figure 16. Create Logical Drive — Introduction window

2. The Introduction (Create Logical Drive) window shows the amount of free or unconfigured capacity available and provides a link for tips and examples on allocating capacity. Click **Next** to continue.
3. The Specify Array (Create Logical Drive) window opens to allow you to define the parameters for an array that you will create for the logical drive. In this window, specify the following array parameters:
 - RAID level: RAID-0, RAID-1, RAID-3 or RAID-5
 - Drive selection choices: Select the number of drives that will be used for this array. You can either specify the number of drives in the array and let the storage subsystem make the selection of which optimal hard drives in the storage subsystem to use, or you can manually select which hard drives will form the array.

Note: In manual mode, to select multiple drives, you can press and hold the **Ctrl** key and then use the mouse to select the desired hard drive or drives in the available unconfigured optimal hard drive list.

Click **Next** to continue.

- A box opens stating that you have successfully set the parameters for creating an array. Click **OK** to continue. The Specify Capacity/Name (Create Logical Drive) window opens.
4. In the Specify Capacity/Name (Create Logical Drive) window, specify the following logical drive parameters for the logical drive you want to create:

Logical drive capacity

The new logical drive capacity can either be the entire unconfigured capacity in an array or a portion of the capacity.

Units You can choose from GB, MB, and TB.

Logical drive name

This name must be unique in the storage subsystem.

Settings

You can create the logical drive either using the DS4000 Storage Server recommended settings or you can select your own customized I/O characteristics, controller ownership, and logical-drive-to-LUN mapping settings.

If you choose to use the DS4000 Storage Server default settings, click the check box next to "Use recommended settings," click **Next** then go to step 6. If you want to select your own customized settings, click the check box next to "Customize settings (I/O characteristics and controller ownership)," click **Next** then go to step 5.

5. The Customize Advanced Logical Drive Parameters (Create Logical Drive) window opens. In the Customize Advanced Logical Drive Parameters (Create Logical Drive) window, specify the following settings:
 - The appropriate I/O characteristics (the I/O characteristics type, segment size, and cache read-ahead multiplier)
The I/O characteristics settings can be manually specified or automatically set based on the selection of one of the following expected logical drive usages: file system, database, or multimedia.
 - The preferred controller ownership

Click **Next**.

6. The Specify Logical Drive-to-LUN Mapping (Create Logical Drive) window opens.

In the Specify Logical Drive-to-LUN Mapping (Create Logical Drive) window, specify how you want the logical drive-to-LUN mapping to be mapped to the individual logical drive.

The logical drive-to-LUN mapping preference can be one of the following two settings:

Automatic

The Automatic setting specifies that a LUN is automatically assigned to the logical drive using the next available LUN within the default host group. This setting grants logical drive access to host groups or host computers that have no specific logical drive-to-LUN mappings (those that were designated by the default host group node in the Topology view). If the Storage Partition feature is not enabled, you must specify the Automatic setting.

Map later using the Mappings View

This setting specifies that you are not going to assign a LUN to the logical drive during creation. This setting enables you to define a specific logical drive-to-LUN mapping and create storage partitions using the Mappings Defined option. When you enable storage partitioning, specify this setting.

7. Click **Finish** to create the logical drive. The Create Logical Drive Wizard – Creation Successful window opens and states that the logical drive was successfully created.
8. If you do not want to create another logical drive, go to step 10. Otherwise, indicate in the Creation Successful (Create Logical Drive) window if you want to create the logical drive from the same array, if there is free capacity, or from a different array if there is no capacity in the array in which the logical drive was just created. Click **Yes** and go to step 9. If you do not want to create another logical drive, click **No** and continue with step 10.
9. The Allocate Capacity (Create Logical Drive) window opens. Select if you want to create the new logical drive from free capacity on the same array, free capacity on a different array, or from unconfigured capacity (create a new array). Then continue with step 4 on page 142.
10. The Completed (Create Logical Drive) window opens. Click **OK**.
11. Register the logical drive with the operating system.

After you create logical drives with automatic logical drive-to-LUN mappings, follow the appropriate instructions in the Installation and Support Guide for your operating system to enable the operating system to discover the new logical drive.

Configuring alert notifications

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

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

For more information about notification options, see the *IBM DS4000 Storage Manager Concepts Guide*.

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

1. Ensure that the installation CD is inserted in the CD-ROM drive on your designated NMS.
2. From the installation CD, copy the SM9.MIB file from the SM9mib directory to the NMS.
3. Follow the steps required by your NMS to compile the MIB. For details, contact your network administrator or see the documentation for the storage management product you are using.

Renaming storage subsystems

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

Creating a storage subsystem profile

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

Complete the following steps to create a storage subsystem profile:

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

If you have a storage server with firmware 04.01.xx.xx or earlier, you must also perform a controller state capture in addition to saving the storage subsystem profile. To use the Capture State Information feature, upgrade your controller to firmware 04.01.02.34 or later.

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

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

Performing other storage subsystem management tasks

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

- Locate a storage subsystem
- View a storage subsystem profile
- Configure a storage subsystem password
- Create and manage arrays and array groups
- Use the performance monitor
- Create storage partitions (if Storage Partitioning is enabled on your system)
- Download drive firmware

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 and configuring storage management software on AIX host systems,” on page 23, Chapter 3, “Installing and configuring storage management

software on HP-UX host systems,” on page 63, Chapter 4, “Installing and configuring storage management software on Solaris host systems,” on page 83, or Chapter 5, “Using DS4000 with Linux on POWER-based hosts,” on page 121 for installation instructions.

Appendix A. MC/Service Guard configuration details

Note: The Hewlett Packard publication *Managing MC/ServiceGuard* suggests using **vgimport -m -s** with LVM commands during the configuration distribution process. It is important to consider, however, that the **vgimport** command does not preserve the primary and alternate paths of the imported volume groups.

When **vgimport** reconstructs the newly imported volume groups, entries are made in the `/etc/lvmtab` file. In this process of reconstruction, the system reorders disks in the file. The revised order of LUNs on the list causes LUNs to remain on non-preferred paths after failover. You might expect this condition to occur during distribution of the volume groups using the following LVM command:

```
vgimport -s -m /tmp/vg_group_name.map /dev/vg_group_name
```

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

This precaution applies only when you use the `-s` option to import devices or disks with redundant paths. The condition occurs because the `-s` option causes a search on the system for each disk when used with **vgimport**.

When the new list is prepared, the links that were previously designated as primary and alternate might not remain as they had been configured on node *A*.

Complete the following steps to correct the paths that were changed in this way when using **vgimport -m -s** with LVM commands.

Note: The following instructions are for a two-node cluster. You can easily expand or adapt these instructions to accommodate a cluster which has more than two nodes.

1. Identify and note the primary and alternate paths for each device by typing the following command:

```
# SMdevices
```

2. Compare the primary and alternate paths for each device on node *A* to Node *B* by using the **vgdisplay -v** command. As previously noted, an expected behavior of the **vgimport -m -s** command is that the primary and alternate paths change during export.

On node *A*:

```
#vgdisplay -v volume_group_name
```

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

On node *B*:

```
#vgdisplay -v volume_group_name
```

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

3. If the original primary path of a disk becomes an alternate path after the newly imported volume group entry is created in `/etc/lvmtab`, the order can be easily reverted by using **vgreduce** to remove the wrong primary path and then using

vgextend to add the path back again:

```
#vgreduce  
vg1 /dev/dsk/device_name#vgextend  
vg1 /dev/dsk/device_name
```


Appendix B. JNI and QLogic host bus adapter settings

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

JNI HBA card settings

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

Table 26. Configuration settings for FCI-1063

Original value	New value	Comments
fca_nport = 0;	fca_nport = 1;	
ip_disable = 0;	ip_disable = 1;	
failover = 30;	failover = 60;	
busy_retry_delay = 5000;	busy_retry_delay = 5000;	
link_recovery_delay = 1000;	link_recovery_delay = 1000;	
scsi_probe_delay = 5000;	scsi_probe_delay = 5000;	
def_hba_binding = "fca-pci*";	def_hba_binding = "nonjni";	
def_wwnn_binding = "\$xxxxxx"	def_wwnn_binding = "xxxxxx"	Uncomment the line. Remove the \$.
def_wwpn_binding = "\$xxxxxx"	Same as the original entry.	Uncomment the line.
Will be added by reconfigure script	name="fca-pci" parent="physical path" unit-address="#"	Uncomment the line.
Will be added by reconfigure script	target0_hba="fca-pci0" target0_wwpn="controller wwpn";	
Will be added by reconfigure script	name="fca-pci" parent="physical path" unit-address="#"	
Will be added by reconfigure script	target0_hba="fca-pci1" target0_wwpn="controller wwpn";	

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

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

Table 27. Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562

Original value	New value	Comments
FcLoopEnabled = 1	FcLoopEnabled = 0	Uncomment the line.
FcFabricEnabled = 0	FcFabricEnabled = 1	Uncomment the line.
FailoverDelay = 30;	FailoverDelay = 60;	Uncomment the line.

Table 27. Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562 (continued)

Original value	New value	Comments
def_hba_binding = JNIC146x	Direct-attached configurations: def_hba_binding = JNIC146x SAN-attached configurations: def_hba_binding = null	Setting depends on whether you have a direct-attached or SAN-attached configuration.
JniCreationDelay = 5;	JniCreationDelay = 10;	Uncomment the line.
def_wwnn_binding = "\$xxxxxxxxxxx"	def_wwnn_binding = "xxxxxxxxxxx"	Uncomment the line.
def_wwpn_binding = "\$xxxxxxxxxxx"	Same as the original entry.	Uncomment the line.
Add	target0_hba = "jnic146x0";	
Add	target0_wwpn = "<controller_wwpn>"	
Add	target1_hba = "jnic146x1";	
Add	target1_wwpn = "<controller_wwpn>"	

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

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

Table 28. Configuration settings for FCE-1063/FCE2-1063/FCE-6410/FCE2-6410

Original value	New value	Comment
FcLoopEnabled = 1	FcLoopEnabled = 0	Uncomment the line.
FcFabricEnabled = 0	FcFabricEnabled = 1	Uncomment the line.
FailoverDelay = 30;	FailoverDelay = 60;	Uncomment the line.
JniCreationDelay = 5;	JniCreationDelay = 10;	Uncomment the line.
def_wwnn_binding = "\$xxxxxxxxxxx"	def_wwnn_binding = "xxxxxxxxxxx"	Uncomment the line.
def_wwpn_binding = "\$xxxxxxxxxxx"	Same as the original entry.	Uncomment the line.
Add	target0_hba = "jnic146x0";	
Add	target0_wwpn = "controller_wwpn"	
Add	target1_hba = "jnic146x1";	
Add	target1_wwpn = "controller_wwpn"	

Note:

- You might need to run the `/etc/raid/bin/genjnicnf` reconfigure script from the Solaris shell.

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

- Set `portEnabled = 1`; only when you see JNI cards entering non-participating mode in the `/var/adm/messages` file. Under that condition:
 - Set `FcPortCfgEnabled = 1`;

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 29. Configuration settings for FC64-1063

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

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

```
# /etc/raid/bin/gen SCSIconf
```

QLogic HBA card settings

The QLogic cards are not plug-and-play with auto-configuration. Instead, you need to change the settings or bindings, as described in Table 30.

Note: In Table 30, the HBA is identified as **hba0**. However, you need to modify the settings on both QLogic HBA cards: **hba0** and **hba1**.

When you modify the settings on **hba1** use the same values that are listed in the table, but change all instances of **hba0** to **hba1**, as shown in the following example:

	Original value	New value
hba0	hba0-execution-throttle=16;	hba0-execution-throttle=255;
hba1	hba1-execution-throttle=16;	hba1-execution-throttle=255;

In the Vi Editor, uncomment and modify the loop attributes of each QLogic HBA card, using the values described in Table 30.

Table 30. Configuration settings for QL2342

Original value	New value	Comments
hba0-max-frame-length=2048;	hba0-max-frame-length=2048	Use the default.
hba0-execution-throttle=16;	hba0-execution-throttle=255;	Change.
hba0-login-retry-count=8;	hba0-login-retry-count=30;	Change.
hba0-enable-adapter-hard-loop-ID=0;	hba0-enable-adapter-hard-loop-ID=1;	Change.
hba0-adapter-hard-loop-ID=0;	hba0-adapter-hard-loop-ID=0;	Needs to be a unique number.
hba0-enable-LIP-reset=0;	hba0-enable-LIP-reset=0;	Use the default.
hba0-enable-LIP-full-login=1;	hba0-enable-LIP-full-login=1;	Use the default.
hba0-enable-target-reset=0;	hba0-enable-target-reset=0;	Use the default.
hba0-reset-delay=5	hba0-reset-delay=8	Change.
hba0-port-down-retry-count=8;	hba0-port-down-retry-count=70;	Change.
hba0-maximum-luns-per-target=8;	hba0-maximum-luns-per-target=0;	Change.
hba0-connection-options=2;	hba0-connection-options=2;	Use the default.
hba0-fc-tape=1;	hba0-fc-tape=0;	Change.

Appendix C. Downloading ESM firmware and drive firmware

The information in this appendix describes how to download current environmental service module (ESM) firmware to an DS4000 EXP storage expansion enclosure and current drive firmware to a DS4000 storage server, using the following procedures:

- “Verifying and downloading ESM firmware”
- “Downloading drive firmware” on page 156

Verifying and downloading ESM firmware

This section describes how to verify current ESM firmware levels and how to download ESM firmware to an DS4000 EXP storage expansion enclosure, using the following procedures:

- “Verifying current ESM firmware version”
- “Downloading ESM firmware” on page 154

Important: If you are using Storage Manager controller firmware version 5.41.xx.xx or 5.42.xx.xx with the EXP100, the EXP100 ESM firmware must be at version 955x or higher. However, you cannot upgrade the EXP100 firmware from 953x to 955x.

Before proceeding, see the *IBM TotalStorage DS4000 EXP100 storage expansion enclosure Installation and User's Guide* or see the readme for controller firmware version 5.42.xx.xx at the following Web site:

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

Click the link for **DS4100 (FAStT100)**. When the page opens, click the Download tab, then click the link for **Current recommended Firmware (v5.42) and Storage Manager (v8.42)**.

Verifying current ESM firmware version

Choose one of the following methods to determine the storage expansion enclosure's ESM firmware version, using DS4000 Storage Manager.

- **Method One:**

1. In the Subsystem Management window, select **Storage Subsystem ► View Profile**. The Storage Subsystem Profile window opens.
2. Select the **Enclosures** tab and scroll through the Profile For Storage Subsystem panel to locate the ESM card firmware version information.

Tip: The Profile for Storage Subsystem panel contains all the profile information for the entire subsystem. Therefore, it might be necessary to scroll through a large amount of information to locate the ESM firmware version information.

- **Method Two:**

1. In the Physical View panel of the Subsystem Management window, click **Advanced → Maintenance → Download → ESM firmware**. The Download Environmental (ESM) Card Firmware window opens.
2. In the Download Environmental (ESM) Card Firmware window, locate the firmware version of each ESM in the storage expansion enclosure.

Tip: Each storage expansion enclosure has two ESM cards: A and B.

Downloading ESM firmware

To find the latest versions of ESM firmware files that you need to download, see the most recent readme file, which is located at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for your operating system and click the link in the Current Version column.

Complete the following procedures to download ESM firmware:

- “Preparing for the download”
- “Downloading the ESM firmware” on page 155

Preparing for the download

To prepare for the download, review and complete all of the steps in this section before you start the download procedures that are listed in “Downloading the ESM firmware” on page 155.

Before you begin: Ensure that you have the latest version of the DS4000 Storage Manager installed on your DS4000 management station. You can download the latest version from the following IBM DS4000 support Web site:

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

Complete the following steps to prepare your system for the firmware download:

1. Backup your logical drives

Complete a full backup of all logical drives that are defined in the DS4000 storage server.

4. Save the storage subsystem profile

Saving a profile is necessary in order to restore a storage subsystem configuration in the event of a catastrophic failure. Save your storage subsystem profile by completing the following steps:

1. In the Subsystem Management Menu, select **View → Storage Subsystem Profile**. The Storage Subsystem Profile window opens.
2. In the Storage Subsystem Profile window, click **Save As**. The Save Profile window opens.
3. In the Save Profile window, select or enter the file directory and file name. Click **Save** to save the profile. The Save Profile window closes.

Attention: Do *not* store DS4000 Storage Server profiles in the same location as the logical drives that are defined on your DS4000 Storage Server. If you do, and the logical drives fail during the subsequent migration, you will lose data.

4. Click **Close** to close the Storage Subsystem Profile window.

5. Verify current EXP ESM status

In the DS4000 Storage Manager Subsystem Management window, verify that all ESMs in the configuration that you are upgrading are in the Optimal state, and there are not any lost drive-path redundancy errors. Correct any problems using the suggested instructions in the DS4000 Storage Manager Recovery Guru window.

To verify, either click the **Recovery Guru** icon in the Subsystem Management window, or click the **Enclosure Status** icon in each storage expansion enclosure with the red X mark overlay, as displayed in the Physical View.

6. Stop all subsystem monitoring operations

Subsystem monitoring operations such as temperature monitoring, fan and power supply status checking, and drive insertion or removal, must be stopped during the ESM firmware download.

Note: Error reporting for failed components is deferred until after the download is complete.

7. Determine whether you need to stop I/O

Downloading firmware without stopping I/O is referred to as *concurrent download*. Concurrent ESM firmware download is supported with the following conditions:

- You can only use concurrent ESM firmware download with Fibre Channel storage expansion enclosures. With SATA storage expansion enclosures, you must stop I/O before downloading ESM firmware.
- You must download the ESM firmware to one storage expansion enclosure at a time. If you want to download ESM firmware to multiple storage expansion enclosures at the same time, you must stop I/O before downloading.
- Ensure that concurrent ESM download is supported with the new ESM firmware package by checking its readme file before you download the package. If concurrent download is not supported with the new firmware, you must stop I/O before downloading the firmware.

When you are done, continue to “Downloading the ESM firmware.”

Downloading the ESM firmware

Before you begin: To find the latest versions of ESM firmware files that you need to download, see the most recent readme file, which is located at the following Web site:

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

Click the link for your storage server.

When the page opens, click the Download tab, then click the link for your Storage Manager firmware level. In the tables, find the Storage Manager listing for your operating system and click the link in the Current Version column.

Complete the following steps to download the ESM firmware to the DS4000 EXP storage expansion enclosure:

1. In the Subsystem Management window, select **Advanced** —> **Maintenance** —> **Download** —> **ESM firmware**. The Download Environmental (ESM) Card Firmware window opens.
2. In the Download Environmental (ESM) Card Firmware window, select the storage expansion enclosure to be updated with new firmware, and then click **Browse** to open the Select Environmental (ESM) Card Firmware File window.

Attention: You can download the ESM firmware to several storage expansion enclosures at the same time by selecting more than one storage

expansion enclosure – however, you must stop all I/O before performing the download. If you do not want to stop I/O, then download the ESM firmware to each storage expansion enclosure, one at a time.

3. In the Select Environmental (ESM) Card Firmware File window, select the directory that contains the ESM firmware file, then select the name of the ESM firmware file. Click **OK** to close the window and continue.

Note: You might need to select All Files (*.*) in the “Files of type” pull-down menu to display the correct file name for selection. If you do, then a warning window opens. Click **OK** when the warning window opens.

4. In the Download Environmental (ESM) Card Firmware window, click **Start**. When the Confirm Download window opens, type yes and click **OK**. The ESM firmware download begins. During the download, the status field of the Download Environmental (ESM) Card Firmware window displays the progress of each storage expansion enclosure.

When the status field of each storage expansion enclosure displays a status of Complete, the ESM firmware download is complete.

Downloading drive firmware

This section provides instructions for downloading DS4000 drive firmware. See the online help for additional information.

Important:

1. IBM supports firmware download with I/O, sometimes referred to as “concurrent firmware download”. This feature is not supported for drive firmware.
2. Before starting the drive firmware download process:
 - Stop all I/O activity before downloading drive firmware to a DS4000 Storage Server.
 - Unmount the file systems on all logical drives accessing the drives that you select for firmware upgrade.
 - Complete a full backup of all data residing on the drives that you select for firmware upgrade.

Downloading Storage Manager version 9.10 drive firmware

To download drive firmware for DS4000 Storage Manager version 9.10, perform the following steps:

1. From the Enterprise Management window, select a storage subsystem.
2. On the Enterprise Management window’s menu bar, click **Tools** → **Manage Device**. The Subsystem Management window opens.
3. On the Subsystem Management window’s menu bar, click **Advanced** → **Maintenance** → **Download** → **Drive Firmware/Mode pages ...**. The Download Drive Firmware wizard window opens to the Introduction page. Read the instructions displayed and click **Next**.

Note: Storage Manager version 9.10 offers you the option to download and update up to four different firmware file types simultaneously. Previous Storage Manager versions did not offer this capability.

4. Click **Browse** to locate the server directory that contains the firmware that you plan to download.

5. Select the firmware file that you plan to download and click **OK**. The file appears listed in the "Selected Packages" window pane.
6. Select the firmware file for any additional drive types that you intend to download and click **OK**. Additional files appear listed in the "Selected Packages" window pane. A maximum total of four drive types are possible.
7. Click **Browse** to repeat step 6 until you have selected each firmware file that you plan to download.
8. When you have finished specifying the firmware packages for download, select **Next**.
9. The Select Drive window opens, containing two tabs, a **Compatible Drives** tab and an **Incompatible Drives** tab. The **Compatible Drives** tab contains a list of the drives compatible to the firmware package types that you selected. From that list, select the drives to which you plan to download the drive firmware that you selected in steps 6 and 7.

Note: The firmware that you propose to download should be listed on the **Compatible Drives** tab. If your particular drives' product ID matches the firmware type, however, and it is not listed as compatible on the tab, contact your IBM technical support representative for additional instructions.

10. Select the **Compatible Drives** tab.
Press and hold the **Ctrl** key while using your mouse to select multiple drives individually, or press and hold the **shift** key while using your mouse to select multiple drives listed in series. The compatible firmware that you selected in steps 5 and 6 will download to the drives that you select.
11. Click **Finish** to initiate download of the drive firmware to each compatible drive that you selected in step 9.
12. The Download Drive Firmware warning opens and prompts: "Do you want to continue?" Type **yes** and click **OK** to start the drive firmware download.
13. The Download Progress window opens. Do not intervene until the download process completes.
14. Every drive scheduled for firmware download will be designated as in progress until successful or failed.

Note: Complete the following two steps if you receive a failure.

- a. Click the **Save as** button to save the error log.
- b. On the Subsystem Management window's menu bar, click **Advanced** → **Trouble Shooting** → **Open Event Log** and complete the following tasks necessary to save the storage subsystem event log before contacting your IBM Service Representative and proceeding to step 16.
 - 1) Click the **Select all** button.
 - 2) Click **Save the Storage Subsystem Event Log**.
15. When the **Close** button appears active, the drive firmware download process is complete.
16. Click **Close** to exit the Download Progress window.
17. Use either of the following procedures to determine or verify what level of drive firmware resides on a particular drive:
 - Right-click on that drive in the Logical/Physical View in the Subsystem Management window and click **Properties**. The associated drive firmware version will be listed in the drive properties table.

- Right-click on **Storage Subsystem** → **View Profile** in the Logical/Physical View of the Subsystem Management window.

Downloading Storage Manager version 8.x drive firmware

To download drive firmware for DS4000 Storage Manager version 8.x, perform the following steps:

1. From the Enterprise Management window, select a storage subsystem.
2. Click **Tools** → **Manage Device**. The Subsystem Management window opens.
3. Click **Advanced** → **Download** → **Drive Firmware**. The Download Drive Firmware window opens. It presents a list of all installed drives.
4. From the list of all installed drives, select the drive to which you plan to download drive firmware.

Press and hold the **Ctrl** key while using your mouse to select multiple drives singly, or press and hold the **shift** key while selecting multiple drives in series by using your mouse pointer.

5. Click **Browse** to identify the directory that contains the drive firmware that you intend to download.
6. Select the file containing the drive firmware that you intend to download.
7. Select **Start**. The Confirm Download window appears.
8. Type **yes** when the Confirm Download window prompts: Are you sure you want to continue?
9. The download begins and the Status window indicates progress. Do not intervene until the download process completes.
10. When the **Cancel** button appears active, the download process is complete. Press **Cancel** to close the Status window.
11. To determine or verify the level of drive firmware residing on a particular drive, click **View** → **Subsystem Profile** in the Logical/Physical View of the Subsystem Management window.

Appendix D. Additional DS4000 documentation

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

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

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

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

DS4500 Fibre Channel Storage Server library

Table 31 associates each document in the DS4500 (previously FAStT900) Fibre Channel Storage Server library with its related common user tasks.

Table 31. TotalStorage DS4500 Fibre Channel Storage Server document titles by user tasks

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

DS4400 Fibre Channel Storage Server library

Table 32 associates each document in the DS4400 (previously FAStT700) Fibre Channel Storage Server library with its related common user tasks.

Table 32. TotalStorage DS4400 Fibre Channel Storage Server document titles by user tasks

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

DS4300 Fibre Channel Storage Server library

Table 33 associates each document in the DS4300 (previously FAStT600) Fibre Channel Storage Server library with its related common user tasks.

Table 33. TotalStorage DS4300 Fibre Channel Storage Server document titles by user tasks

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

DS4100 SATA Storage Server library

Table 34 associates each document in the DS4100 (previously FAStT100) SATA Storage Server library with its related common user tasks.

Table 34. TotalStorage DS4100 SATA Storage Server document titles by user tasks

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

Other DS4000 and DS4000-related documents

Table 35 associates each of the following documents with its related common user tasks.

Table 35. TotalStorage DS4000 and DS4000–related document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
<i>IBM Safety Information, P48P9741</i>					✓	
<i>IBM TotalStorage DS4000 Quick Start Guide, GC26-7738</i>	✓	✓				
<i>IBM TotalStorage DS4000 Hardware Maintenance Manual,GC26-7702</i>						✓
<i>IBM TotalStorage DS4000 Problem Determination Guide, GC26-7703</i>						✓
<i>IBM Fibre Channel Planning and Integration: User's Guide and Service Information, SC23-4329</i>	✓	✓			✓	✓
<i>IBM TotalStorage DS4000 FC2-133 Host Bus Adapter Installation and User's Guide, GC26-7736</i>		✓			✓	
<i>IBM TotalStorage DS4000 FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7737</i>		✓			✓	
<i>IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview GC26-7713</i>	✓	✓	✓	✓		
<i>Fibre Channel Solutions - IBM DS4000 EXP500 Installation and User's Guide, 59p5637</i>	✓	✓		✓	✓	✓

Table 35. TotalStorage DS4000 and DS4000–related document titles by user tasks (continued)

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

Appendix E. Connecting HBAs in a FC switch environment

When connecting Fibre Channel host bus adapters in host servers to DS4000 Storage Server host ports in a Fibre Channel switch environment, IBM recommends that you establish one-to-one zones between the HBA and DS4000 host ports.

For instructions on how to set up a zone inside a Fibre Channel switch, see your Fibre Channel switch documentation. You can find links to switch documentation at the following Web site:

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

Figure 17 illustrates a sample configuration consisting of two host servers, each with two Fibre Channel host bus adapters, each of which are connected to a DS4000 Storage Server host through a Fibre Channel switch.

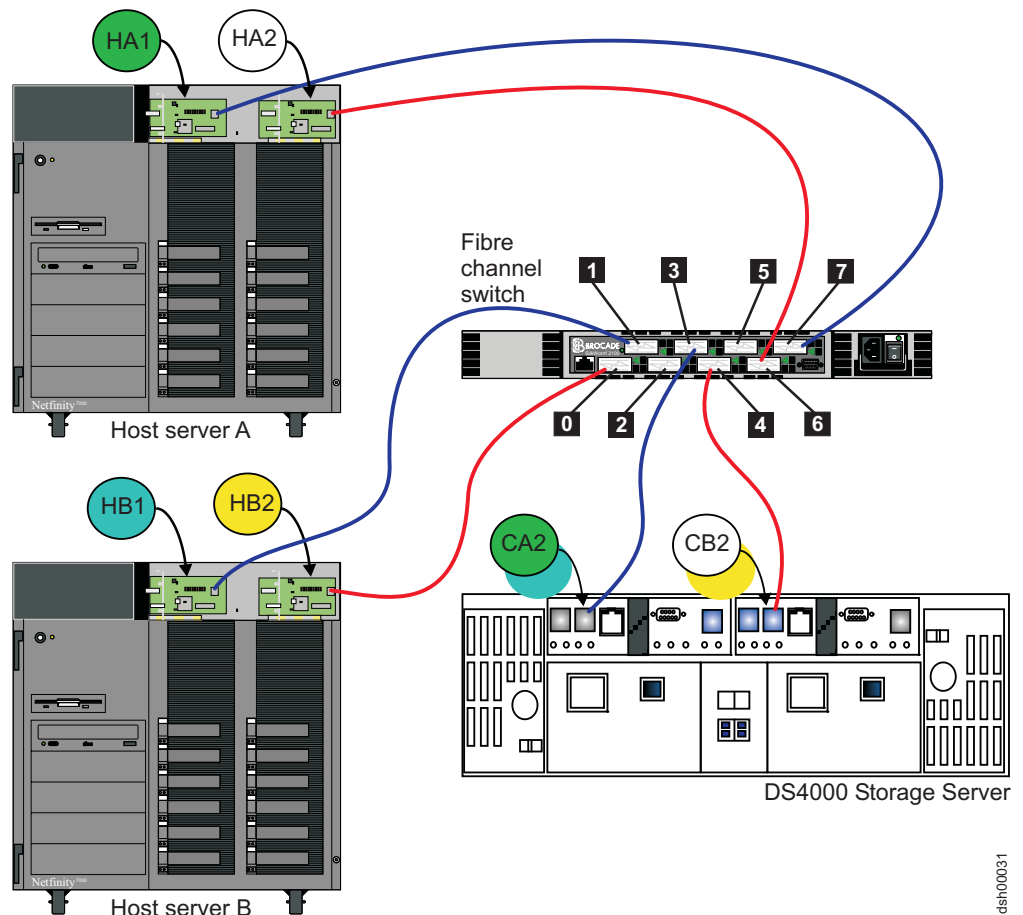


Figure 17. Fibre Channel fabric switch environment connecting host servers to a DS4000 Storage Server

Legend:

Controller	Zone	Port	Connection to and from	Description
A	1	7 3	HA1 CA2	This will zone the HBA port HA1 in host server A to port 2 of the DS4300 Storage Server controller A.
	2	1 3	HB1 CA2	This will zone the HBA port HB1 in host server B to port 2 of the DS4300 Storage Server controller A.
B	3	6 4	HA2 CB2	This will zone the HBA port HA2 in host server A to port 2 of the DS4300 Storage Server controller B.
	4	0 4	HB2 CB2	This will zone the HBA port HB2 in host server B to port 2 of the DS4300 Storage Server controller B.

Appendix F. Accessibility

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

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

Table 36 defines the keyboard operations that enable you to navigate, select, or activate user interface components. The following terms are used in the table:

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

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

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

Table 36. DS4000 Storage Manager alternate keyboard operations

Short cut	Action
F1	Open the Help.
F10	Move keyboard focus to main menu bar and post first menu; use the arrow keys to navigate through the available options.
Alt+F4	Close the management window.
Alt+F6	Move keyboard focus between dialogs (non-modal) and between management windows.
Alt+ underlined letter	Access menu items, buttons, and other interface components by using the keys associated with the underlined letters. For the menu options, select the Alt + underlined letter combination to access a main menu, and then select the underlined letter to access the individual menu item. For other interface components, use the Alt + underlined letter combination.
Ctrl+F1	Display or conceal a tool tip when keyboard focus is on the toolbar.
Spacebar	Select an item or activate a hyperlink.

Table 36. DS4000 Storage Manager alternate keyboard operations (continued)

Short cut	Action
Ctrl+Spacebar (Contiguous/Non-contiguous) AMW Logical/Physical View	<p>Select multiple drives in the Physical View.</p> <p>To select multiple drives, select one drive by pressing Spacebar, and then press Tab to switch focus to the next drive you want to select; press Ctrl+Spacebar to select the drive.</p> <p>If you press Spacebar alone when multiple drives are selected then all selections are removed.</p> <p>Use the Ctrl+Spacebar combination to deselect a drive when multiple drives are selected.</p> <p>This behavior is the same for contiguous and non-contiguous selection of drives.</p>
End, Page Down	Move keyboard focus to the last item in the list.
Esc	Close the current dialog (does not require keyboard focus).
Home, Page Up	Move keyboard focus to the first item in the list.
Shift+Tab	Move keyboard focus through components in the reverse direction.
Ctrl+Tab	Move keyboard focus from a table to the next user interface component.
Tab	Navigate keyboard focus between components or select a hyperlink.
Down arrow	Move keyboard focus down one item in the list.
Left arrow	Move keyboard focus to the left.
Right arrow	Move keyboard focus to the right.
Up arrow	Move keyboard focus up one item in the list.

Notices

This information 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 grant 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.*

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publications. IBM may make improvements or changes (or both) in the product(s) or program(s) (or both), described in this publication at any time without notice.

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

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance,

compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

Some software may differ from its retail version (if available), and may not include user manuals or all program functionality.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM
AIX
@server
FlashCopy
Netfinity
POWER
pSeries
RS/6000
TotalStorage

Microsoft Windows and Windows NT are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

This glossary provides definitions for the terminology and abbreviations used in IBM TotalStorage DS4000 publications.

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

www.ibm.com/ibm/terminology

This glossary also includes terms and definitions from:

- *Information Technology Vocabulary* by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.
- *IBM Glossary of Computing Terms*. New York: McGraw-Hill, 1994.

The following cross-reference conventions are used in this glossary:

See Refers you to (a) a term that is the expanded form of an abbreviation or acronym, or (b) a synonym or more preferred term.

See also

Refers you to a related term.

Abstract Windowing Toolkit (AWT). A Java graphical user interface (GUI).

accelerated graphics port (AGP). A bus specification that gives low-cost 3D graphics cards faster access to main memory on personal computers than the usual peripheral component interconnect (PCI) bus. AGP reduces the overall cost of creating high-end graphics subsystems by using existing system memory.

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

adapter. A printed circuit assembly that transmits user data input/output (I/O) between the internal bus of the

host system and the external fibre-channel (FC) link and vice versa. Also called an I/O adapter, host adapter, or FC adapter.

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

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

AGP. See *accelerated graphics port*.

AL_PA. See *arbitrated loop physical address*.

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

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

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

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

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

ATA. See *AT-attached*.

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

auto-volume transfer/auto-disk transfer (AVT/ADT).

A function that provides automatic failover in case of controller failure on a storage subsystem.

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

AWT. See *Abstract Windowing Toolkit*.

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

BIOS. See *basic input/output system*.

BOOTP. See *bootstrap protocol*.

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

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

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

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

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

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

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

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

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

CRC. See *cyclic redundancy check*.

CRT. See *cathode ray tube*.

CRU. See *customer replaceable unit*.

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

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

dac. See *disk array controller*.

dar. See *disk array router*.

DASD. See *direct access storage device*.

data striping. See *striping*.

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

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

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

DHCP. See *Dynamic Host Configuration Protocol*.

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

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

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

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

DMA. See *direct memory access*.

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

DRAM. See *dynamic random access memory*.

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

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

ECC. See *error correction coding*.

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

EISA. See *Extended Industry Standard Architecture*.

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

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

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

E_port. See *expansion port*.

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

ESD. See *electrostatic discharge*.

ESM canister. See *environmental service module canister*.

EXP. See *storage expansion enclosure*.

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

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

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

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

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

FAST MSJ. See *FAST Management Suite Java*.

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

FC. See *fibre channel*.

FC-AL. See *arbitrated loop*.

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

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

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

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

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

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

F_port. See *fabric port*.

FRU. See *field replaceable unit*.

GBIC. See *gigabit interface converter*

gigabit interface converter (GBIC). A transceiver that performs serial, optical-to-electrical, and

electrical-to-optical signal conversions for high-speed networking. A GBIC can be hot swapped. See also *small form-factor pluggable*.

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

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

graphical user interface (GUI). A type of computer interface that presents a visual metaphor of a real-world scene, often of a desktop, by combining high-resolution graphics, pointing devices, menu bars and other menus, overlapping windows, icons, and the object-action relationship.

GUI. See *graphical user interface*.

HBA. See *host bus adapter*.

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

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

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

host computer. See *host*.

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

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

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

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

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

IC. See *integrated circuit*.

IDE. See *integrated drive electronics*.

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

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

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

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

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

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

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

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

IP. See *Internet Protocol*.

IPL. See *initial program load*.

IRQ. See *interrupt request*.

ISA. See *Industry Standard Architecture*.

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

JRE. See *Java Runtime Environment*.

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

LAN. See *local area network*.

LBA. See *logical block address*.

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

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

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

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

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

loop group. A collection of storage area network (SAN) devices that are interconnected serially in a single loop circuit.

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

LPAR. See *logical partition*.

LUN. See *logical unit number*.

MAC. See *medium access control*.

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

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

MCA. See *micro channel architecture*.

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

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

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

MIB. See *management information base*.

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

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

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

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

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

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

MSCS. See *Microsoft Cluster Server*.

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

NMI. See *non-maskable interrupt*.

NMS. See *network management station*.

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

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

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

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

N_port. See *node port*.

NVS. See *nonvolatile storage*.

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

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

ODM. See *Object Data Manager*.

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

partitioning. See *storage partition*.

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

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

PDF. See *portable document format*.

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

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

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

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

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

premium feature key. A file that the storage subsystem controller uses to enable an authorized premium feature. The file contains the feature enable identifier of the storage subsystem for which the

premium feature is authorized, and data about the premium feature. See also *feature enable identifier*.

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

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

PTF. See *program temporary fix*.

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

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

RAID set. See *array*.

RAM. See *random-access memory*.

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

RDAC. See *redundant disk array controller*.

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

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

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

redundant disk array controller (RDAC). (1) In hardware, a redundant set of controllers (either active/passive or active/active). (2) In software, a layer that manages the input/output (I/O) through the active controller during normal operation and transparently reroutes I/Os to the other controller in the redundant set if a controller or I/O path fails.

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

ROM. See *read-only memory*.

router. A computer that determines the path of network traffic flow. The path selection is made from several paths based on information obtained from specific protocols, algorithms that attempt to identify the shortest or best path, and other criteria such as metrics or protocol-specific destination addresses.

RVSD. See *recoverable virtual shared disk*.

SAI. See *Storage Array Identifier*.

SA Identifier. See *Storage Array Identifier*.

SAN. See *storage area network*.

SATA. See *serial ATA*.

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

SCSI. See *small computer system interface*.

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

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

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

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

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

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

SFP. See *small form-factor pluggable*.

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

SL_port. See *segmented loop port*.

SMagent. The DS4000 Storage Manager optional Java-based host-agent software, which can be used on Microsoft Windows, Novell NetWare, HP-UX, and Solaris host systems to manage storage subsystems through the host fibre-channel connection.

SMclient. The DS4000 Storage Manager client software, which is a Java-based graphical user interface (GUI) that is used to configure, manage, and troubleshoot storage servers and storage expansion enclosures in a DS4000 storage subsystem. SMclient can be used on a host system or on a storage management station.

SMruntime. A Java compiler for the SMclient.

SMutil. The DS4000 Storage Manager utility software that is used on Microsoft Windows, HP-UX, and Solaris host systems to register and map new logical drives to the operating system. In Microsoft Windows, it also contains a utility to flush the cached data of the operating system for a particular drive before creating a FlashCopy.

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

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

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

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

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

SRAM. See *static random access memory*.

SSA. See *serial storage architecture*.

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

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

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

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

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

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

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

storage partition topology. In the DS4000 Storage Manager client, the Topology view of the Mappings window displays the default host group, the defined host group, the host computer, and host-port nodes. The host port, host computer, and host group topological elements must be defined to grant access to host computers and host groups using logical drive-to-LUN mappings.

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

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

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

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

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

switch zoning. See *zoning*.

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

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

system name. Device name assigned by the vendor's third-party software.

TCP. See *Transmission Control Protocol*.

TCP/IP. See *Transmission Control Protocol/Internet Protocol*.

terminate and stay resident program (TSR program). A program that installs part of itself as an extension of DOS when it is executed.

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

TL_port. See *translated loop port*.

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

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

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

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

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

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

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

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

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

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

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

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

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

WWN. See *worldwide name*.

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

Index

A

- about this document xi
- access volume/logical drive 4, 10, 26, 51, 70, 90, 125
- additional storage
 - See storage expansion enclosures
- address
 - for e-mail comments xvii
 - for reader comment form xvii
 - of IBM xvii
 - of the IBM director of licensing 169
- agent software
 - AIX 28
 - HP-UX 68
 - package info 7
 - Solaris 89
- AIX
 - boot restrictions 26
 - client software 27
 - concurrent firmware download 30
 - configuration of storage subsystems 29
 - direct-attached configuration 35
 - dynamic capacity expansion (DCE) 57
 - dynamic volume expansion (DVE) 57
 - fast I/O failure 56
 - FCP disk array errors 59
 - fget_config command 43
 - firmware
 - requirements 23
 - upgrading 32
 - hardware requirements 23, 25
 - identifying
 - controller ID numbers 40
 - device names and bus numbers 41
 - FlashCopy logical drive types 47
 - logical drives by operating system device names 46
 - logical drives, redistributing in case of failure 61
 - lsattr command 45
 - lsdev command 42
 - NVSRAM
 - requirements 23
 - upgrading 32
 - operating system requirements 23
 - RDAC driver
 - configuring 38
 - installing 36, 38
 - requirements 36
 - requirements
 - client software 27
 - firmware 23
 - hardware 23, 25
 - operating system 23
 - RDAC driver 36
 - software 25
 - system 25
 - restrictions 26
 - SAN-attached configuration 35

- AIX (*continued*)
 - SMagent 28
 - SMclient 28
 - SMruntime 27
 - software requirements 25
 - staged controller firmware download 31
 - system requirements 25
 - verifying configuration 40
- alert notifications, configuring 143
- appware
 - AIX firmware 23
 - HP-UX firmware 63
 - Linux on POWER firmware 122
 - Solaris firmware 83
- attributes
 - dac 54, 55
 - dar 54
 - hdisk 53, 54, 55
 - LUN 53, 54, 55
 - RDAC driver 38
- audience of this document xi
- Auto Volume Transfer (AVT) 81, 118

B

- boot, AIX restrictions 26
- BOOTP server
 - sample network 12
- bus numbers and device names, identifying 41

C

- cabling instructions xii
- cache mirroring 55, 61
- capture state information feature 144
- catastrophic failure 144
- check drive firmware level 157, 158
- client software
 - AIX 27, 28
 - HP-UX 66, 67
 - package info 6, 7
 - Solaris 86, 87
- cluster services
 - configurations 139
 - MC/Service Guard 147
- comments about this document, how to send xvii
- complete SM SW installation xii
- Concepts Guide xii, xiv
- concurrent download
 - controller firmware 30
 - ESM firmware 155
- configuration
 - details, MC Service Guard 147
 - recovery 144
- configure storage hardware xii
- configure storage subsystems on host xii

- configuring
 - alert notifications 143
 - direct-attached 35, 75, 95, 132
 - RDAC driver 38
 - SAN-attached 35, 75, 95, 132
 - Solaris with JNI host bus adapter cards 96, 97
 - Solaris with QLogic host bus adapter cards 96
 - storage subsystems
 - AIX 29, 35
 - HP-UX 69, 75
 - initial automatic discovery 139
 - Linux on POWER 125, 132
 - Solaris 90, 95
- connect power xii
- controller state capture 144
- controllers
 - firmware 1, 70, 91
 - ID numbers, identifying 40
 - IP addresses 17
- copy services
 - enabling premium features 21
 - overview of premium features 20
- Copy Services Guide xiv

D

- dac (disk array controller)
 - and fget_config command 43
 - and lsattr command 45
 - and RDAC 38
 - attributes 54, 55
 - identifying controller ID numbers 40
 - showing status and location 42
- daisy-chain configuration 35
- dar (disk array router)
 - and fget_config command 43
 - and lsattr command 45
 - and RDAC 38
 - attributes 54
 - showing status 42
 - two dars on one subsystem, causes 39
- DCE (dynamic capacity expansion) 57
- default host type, changing
 - AIX 33
 - HP-UX 74
 - Linux on POWER 129
 - Solaris 94
- determine management method xii
- device drivers
 - DMP
 - configuring VERITAS Volume Manager for 113
 - DMP, installing 107
 - downloading latest versions 2
 - RDAC
 - AIX, file sets 36
 - description 7
 - installing 36, 105
 - viewing attributes 38
- devices
 - identification 39
 - names and bus numbers, identifying 41

- DHCP server
 - sample network 12
- direct management method
 - See out-of-band (direct) management method
- direct-attached configuration 35, 75, 95, 132
- disk array controller
 - See dac
- disk array router
 - See dar
- disk configuration information 144
- DMP
 - See Dynamic Multipathing (DMP)
- documentation
 - documents 2
 - DS4000 xiv, 159
 - DS4000 Storage Manager script xiv
 - files 2
 - online help 4, 140, 144
 - related 163
 - Web sites xvi
- drive firmware
 - parallel download 5
 - using online help 5
- drive firmware download 156
- drive firmware, level 157, 158
- drivers
 - See device drivers
- DS4000 documentation xiv, 159
- DS4000 Storage Manager
 - related documents 163
- DS4000 Storage Manager 9.1
 - user tasks xiv
- DS4000/FASTT product renaming xi
- DS4100 SATA Storage Server library 162
- DS4300
 - restrictions 3, 18
 - usage notes 3, 18
- DS4300 Fibre Channel Storage Server library 161
- DS4400 Storage Server library 160
- DS4500 Fibre Channel Storage Server library 159
- DVE (dynamic volume expansion) 57
- dynamic capacity expansion (DCE) 57
- Dynamic Multipathing (DMP)
 - configuring VERITAS Volume Manager for 113
 - description 7
 - installing 107
 - installing the SMibmasl package 113
 - installing VERITAS Volume Manager packages 110
 - preparing for installation 108
 - system requirements 107
- dynamic tracking 57
- dynamic volume expansion (DVE) 57

E

- e-mail address xvii, 143
- edition notice ii
- enabling
 - multipath I/O with PV-links 75
 - premium features 21
- Enhanced Remote Mirror Option 4, 20

- Enterprise Management window
 - description 7
 - online help 4
- errors, FCP disk array 59
- ESM firmware download 153
- Ethernet
 - and out-of-band (direct) management 10
 - managing subsystems through 10
 - setting up subsystems with 29, 70, 125
 - Solaris requirements 86

F

- fabric switch environment 165
- failure support
 - DMP driver 7
 - RDAC driver 7
 - redistributing logical drives 61, 81, 118
- fast I/O failure 56
- FASTT/DS4000 product renaming xi
- FC cabling instructions xii
- FC HW maintenance and problem determination guide xii
- FC planning and integration user's guide and service information xii
- FC storage server installation guide xii
- FCP disk array errors 59
- fcsci attributes 56, 57
- feature activation code 21
- feature enable identifier 21
- feedback about this document, sending to IBM xvii
- fget_config command 43
- file sets 36
- firmware
 - AIX
 - concurrent firmware download 30
 - downloading latest versions 29
 - requirements 23
 - staged controller firmware download 31
 - upgrading 32
 - description 1
 - drive 156
 - drive firmware download
 - parallel 5
 - using online help 5
 - ESM 153
 - HP-UX
 - downloading latest versions 71
 - requirements 63
 - staged controller firmware download 71
 - upgrading 72
 - Linux on POWER
 - downloading latest versions 126
 - requirements 122
 - staged controller firmware download 126
 - upgrading 127
 - restrictions 3, 7, 8, 18, 19, 20, 28, 30, 31, 68, 72, 89, 92, 153
 - Solaris
 - downloading latest versions 91
 - requirements 83

- firmware (*continued*)
 - Solaris (*continued*)
 - staged controller firmware download 91
 - upgrading 92
- FlashCopy
 - disk array error messages (AIX) 60
 - identifying logical drive types 47
 - overview 20

G

- glossary 171

H

- hardware installation process xii
- hardware maintenance and problem determination xii
- hardware requirements
 - AIX 23, 25
 - HP-UX 63, 65
 - Linux on POWER 122
 - Solaris 83, 85
 - SUSE Linux 124
- HBA
 - See host bus adapters
- HBA installation and user's guide xii
- hdisk
 - and fget_config command 43
 - and lsattr command 45
 - and RDAC 38
 - attributes 53, 54, 55
 - determining preferred path 42
 - setting queue depth 53
 - showing status and location 42
 - verification 39
- heterogeneous environment 139
- host bus adapter installation and user's guide xii
- host bus adapters
 - AIX
 - defining host groups 33
 - displaying a list 34
 - in a direct-attached configuration 35
 - in a SAN-attached configuration 35
 - setting host ports 33
 - single HBA configuration 35
 - HP-UX
 - defining host groups 73
 - displaying a list 73
 - in a direct-attached configuration 75
 - in a SAN-attached configuration 75
 - setting host ports 73
 - Linux on POWER
 - defining host groups 129
 - displaying a list 129
 - in a direct-attached configuration 132
 - in a SAN-attached configuration 132
 - setting host ports 128
 - single HBA configuration 132
 - Solaris
 - in a direct-attached configuration 95
 - in a SAN-attached configuration 95

- host bus adapters (*continued*)
 - Solaris (*continued*)
 - JNI settings 96, 97, 149
 - QLogic settings 104, 152
 - setting host ports 93
 - with switch environment 165
- host group, defining
 - AIX 33
 - HP-UX 73
 - Linux on POWER 128
 - Solaris 93
- host software
 - AIX 28
 - HP-UX 68, 69
 - package info 7
 - Solaris 88, 89
- host type, changing default
 - AIX 33
 - HP-UX 74
 - Linux on POWER 129
 - Solaris 94
- host-agent management method
 - See in-band (host-agent) management method
- how to send your comments xvii
- HP-UX
 - client software 66
 - configuration of storage subsystems 69
 - direct-attached configuration 75
 - firmware
 - requirements 63
 - upgrading 72
 - hardware requirements 63, 65
 - host software
 - installing 68
 - requirements 68
 - logical drives, redistributing in case of failure 81
 - NVSRAM
 - requirements 63
 - upgrading 73
 - PV-links 75
 - requirements
 - client software 66
 - firmware 63
 - hardware 63, 65
 - host software 68
 - NVSRAM 63
 - software 65
 - system 65
 - SAN-attached configuration 75
 - SMagent 68
 - SMclient 67
 - SMruntime 66
 - SMutil 69
 - software requirements 65
 - staged controller firmware download 71
 - system requirements 65
- HyperTerminal version 17

I

- IBM
 - address xvii
 - director of licensing address 169
 - e-mail address xvii, 143
- in-band (host-agent) management method
 - advantages 9
 - description 8
 - diagram 12
 - disadvantages 9
 - utm device 38
 - with SMagent software 7, 28, 68, 89
- install and verify SM SW on host and workstation xii
- install network hardware xii
- install storage expansion enclosure xii
- install storage server/RAID enclosures in rack xii
- installation and support OS guides xii
- installation guide, storage server xii
- installation planning xii
- installation process xii
- installing
 - AIX
 - client software 27, 28
 - RDAC driver 36
 - client software
 - AIX 27
 - HP-UX 66
 - SMclient 28, 67, 87
 - SMruntime 27, 66, 86
 - Solaris 86
 - host software
 - HP-UX 68
 - SMagent 28, 68, 89
 - SMutil 69, 90
 - Solaris 88
 - HP-UX
 - client software 66, 67
 - host software 68
- process diagram for 13
- sequence of 7
- Solaris
 - client software 86, 87
 - host software 88, 89, 90
 - RDAC driver 105

- interoperability
- AIX restrictions 27
- Linux on POWER restrictions 125
- IP addresses for DS4000 controllers 17

J

- JNI
 - cards 96, 97
 - installing driver 96
 - modifying HBA settings 97, 99
 - settings 149

L

- Linux (SUSE)
 - requirements
 - hardware 124
 - software 124
 - system 124
- Linux on POWER
 - configuration of storage subsystems 125
 - direct-attached configuration 132
 - firmware
 - requirements 122
 - upgrading 127
 - hardware requirements 122
 - NVSRAM
 - requirements 122
 - upgrading 128
 - requirements
 - firmware 122
 - hardware 122
 - operating system 121
 - restrictions 124
 - SAN-attached configuration 132
 - staged controller firmware download 126
- load balancing 49
- load_balancing attribute 49
- logical drives
 - creating from free or unconfigured capacity 141
 - expected usage 143
 - redistributing in case of failure 61, 81, 118
 - wizard 141
- loop settings 97, 99
- lsattr command 45
- lsdev command 42
- LUNs
 - and access volumes 10
 - and fget_config command 43
 - attributes 53, 54, 55
 - mapping to a partition
 - AIX 34
 - HP-UX 74
 - Linux on POWER 131
 - Solaris 94
 - maximum number per partition 18

M

- make FC connections xii
- management method
 - in-band (host-agent) 8
 - out-of-band (direct) 10
- management station 3, 5
 - description 3
 - for Linux on POWER 121
- mapping
 - host group
 - AIX 33
 - HP-UX 73
 - Linux on POWER 129
 - Solaris 93

- mapping (*continued*)
 - LUNs, to a partition
 - AIX 34
 - HP-UX 74
 - Linux on POWER 131
 - Solaris 94
 - SMutil 7
- MC/Service Guard 147
- Migration Guide xiv
- minihubs 26, 35, 75, 95, 124, 132
- multipath
 - and AIX fast I/O failure 56
 - DMP, installing on Solaris 107
 - PV-links, using on HP-UX 75
- RDAC
 - AIX devices 38
 - installing on AIX 36
 - installing on Solaris 105
 - Storage Manager software package 6
- redistributing logical drives
 - AIX 61
 - HP-UX 81
 - Solaris 118

N

- new features in this edition 4
- notice, edition ii
- notices
 - general 169
- NVSRAM
 - AIX
 - concurrent download 30
 - requirements 23, 24
 - upgrading 32
 - downloading latest versions 29, 71, 91, 126
- HP-UX
 - requirements 63, 64
 - upgrading 73
- Linux on POWER
 - requirements 122
 - upgrading 128
- Solaris
 - requirements 83, 84
 - upgrading 93
- SUSE Linux
 - requirements 122

O

- Object Data Manager (ODM) attributes
 - definitions 49, 50
 - initial device identification 39
 - lsattr command 54
 - switch_retries and concurrent download 30
 - viewing and setting 48
- online help xii, 4, 140, 144
- operating system
 - AIX
 - requirements 23, 25

- operating system (*continued*)
 - HP-UX
 - requirements 63, 65
 - Linux on POWER
 - requirements 121
 - Solaris
 - requirements 83, 85
 - SUSE Linux
 - requirements 124
- out-of-band (direct) management method
 - advantages 11
 - description 10
 - diagram 12
 - disadvantages 11
 - process diagram 13
 - setting IP addresses 17
- overview of heterogeneous hosts 139

P

- partitioning
 - See Storage Partitioning
- performance
 - in-band (host-agent) management disadvantage 9
 - ODM attribute settings and 53
 - out-of-band (direct) management advantage 11
 - planning 18
- premium features
 - copy services 20
 - enabling 21
 - Storage Partitioning
 - description 21
 - host group 33, 73, 93, 128
 - Web site 21
- prepare for SM SW installation xii
- prepare network connection xii
- prerequisites
 - AIX
 - client software 27
 - firmware 23
 - hardware 23, 25
 - RDAC driver 36
 - software 25
 - system 25
 - HP-UX
 - client software 66
 - firmware 63
 - hardware 63, 65
 - host software 68
 - NVSRAM 63
 - software 65
 - system 65
 - Linux on POWER
 - firmware 122
 - hardware 122
 - Solaris
 - client software 86
 - firmware 83
 - hardware 83, 85
 - host software 88
 - NVSRAM 83

- prerequisites (*continued*)
 - Solaris (*continued*)
 - RDAC driver 88
 - software 85
 - system 85
 - SUSE Linux
 - hardware 124
 - software 124
 - system 124
- products, developed 169
- profile, storage subsystem
 - saving 144
 - viewing 29, 71, 91, 126
- PV-links 75

Q

- QLogic
 - cards 96
 - installing adapter driver 102
 - settings 152
- queue depth for hdisk devices 53

R

- RAID controller enclosure unit installation guide xii
- RDAC driver
 - AIX
 - attributes 38
 - configuring 38
 - installing 36, 38
 - requirements 36
 - description 7
 - Solaris
 - installing 105
 - requirements 88
- reader comment form xvii
- readme files, finding online 2
- recover configuration 144
- redistributing logical drives in case of failure
 - AIX 61
 - HP-UX 81
 - Solaris 118
- Remote Mirror Option 4, 20
- renaming xi
- renaming storage subsystem 144
- requirements
 - AIX
 - client software 27
 - firmware 23
 - hardware 23, 25
 - RDAC driver 36
 - software 25
 - system 25
 - HP-UX
 - client software 66
 - firmware 63
 - hardware 63, 65
 - host software 68
 - NVSRAM 63
 - software 65

requirements (continued)

HP-UX (continued)

system 65

Linux on POWER

firmware 122

hardware 122

Solaris

client software 86

firmware 83

hardware 83, 85

host software 88

NVSRAM 83

RDAC driver 88

software 85

system 85

SUSE Linux

hardware 124

software 124

system 124

resources

documents 2

e-mail address xvii, 143

files 2

online help 4, 140, 144

Web sites xvi, 21

restrictions

AIX 26

DS4100 20

DS4300 3, 18

EXP100 153

firmware 3, 7, 8, 18, 19, 20, 28, 30, 31, 68, 72, 89,
92, 153

in-band (host-agent) management 8

Linux on POWER 124

runtime software

package info 6

S

SAN-attached configuration 26, 35, 75, 95, 124, 132

SATA 7, 24, 64, 84, 123

sending your comments to IBM xvii

services offered in the U.S.A. 169

set link speed xii

setting

IP addresses 17

NMS for SNMP notification 143

Simple Network Management Protocol (SNMP)

notification 143

Simple Network Management Protocol (SNMP)

traps 12

SM concepts guide xii

SM installation and support guide for Linux xii

SM installation and support guide for Novell

NetWare xii

SM installation and support guide for Windows

2000 xii

SMagent

description 7

installing

AIX 28

SMagent (continued)

installing (continued)

HP-UX 68

Solaris 89

software installation sequence 7

SMclient

components 7

description 6

Enterprise Management window 7

installing

AIX 28

HP-UX 67

Solaris 87

software installation sequence 7

Subsystem Management window 7

SMruntime

description 6

installing

AIX 27

HP-UX 66

Solaris 86

software installation sequence 7

SMutil

description 7

installing

HP-UX 69

Solaris 90

software installation sequence 7

snapware

AIX firmware 23

HP-UX firmware 63

Linux on POWER firmware 122

Solaris firmware 83

SNMP

notification 143

traps 12

software installation process xii

software package

RDAC 7

SMagent 7

SMclient 6

SMruntime 6

SMutil 7

software requirements

AIX 25

HP-UX 65

Solaris 85

SUSE Linux 124

Solaris

client software 86

configuration of storage subsystem 90

direct-attached configuration 95

DMP 107, 110, 113, 117

firmware

requirements 83

upgrading 92

hardware requirements 83, 85

host software 88

JNI host bus adapter cards 96, 97

logical drives, redistributing in case of failure 118

- Solaris *(continued)*
 - NVSRAM
 - requirements 83
 - upgrading 93
 - QLogic host bus adapter cards 96
 - RDAC driver
 - installing 105
 - requirements 88
 - requirements
 - client software 86
 - firmware 83
 - hardware 83, 85
 - host software 88
 - NVSRAM 83
 - RDAC driver 88
 - software 85
 - system 85
 - SAN-attached configuration 95
 - SMagent, installing 89
 - SMclient, installing 87
 - SMruntime, installing 86
 - SMutil, installing 90
 - software requirements 85
 - staged controller firmware download 91
 - system requirements 85
 - VERITAS 107, 110, 113, 117
- staged controller firmware download
 - AIX 31
 - HP-UX 71
 - Linux on POWER 126
 - Solaris 91
- starpubs xvii
- start server xii
- storage area network (SAN)
 - configuration
 - AIX 35
 - HP-UX 75
 - Linux on POWER 132
 - Solaris 95
 - technical support Web site xvii
- storage expansion enclosure installation guide xii
- storage expansion enclosures
 - EXP100 7, 23, 24, 64, 84, 122, 123
 - EXP500 23, 63, 83, 122
 - EXP700 23, 63, 83, 122
- storage management station 3
 - description 3, 121
- Storage Manager software
 - installation sequence 7
 - introduction 3
 - list of software packages 5
 - new features 4
 - not supported on Linux on POWER-based hosts 121
 - removing previous versions 27
 - where to obtain 3
- Storage Partitioning
 - AIX restrictions 26
 - and host groups
 - AIX 33
 - HP-UX 73
- Storage Partitioning *(continued)*
 - and host groups *(continued)*
 - Linux on POWER 128
 - Solaris 93
 - configuring storage partitions 142
 - Linux on POWER restrictions 125
 - max. number of LUNs per partition 18
 - overview 21
- storage server installation guide xii
- storage subsystem
 - adding to SMclient
 - AIX 29
 - HP-UX 70
 - Linux on POWER 125
 - Solaris 91
 - configuring
 - AIX 29
 - HP-UX 69
 - Linux on POWER 125
 - procedures 139
 - Solaris 90
 - initial automatic discovery 139
 - introduction 3
 - logical drives
 - identifying by operating system device names 46
 - management methods
 - example showing both types 12
 - in-band (host-agent) 8
 - out-of-band (direct) 10
 - management tasks 144
 - profile, saving 144
 - renaming 144
- Subsystem Management window
 - description 7
 - online help 4
- switch
 - in a SAN-attached configuration 35, 75, 95, 132
 - modifying JNI HBA settings 97, 99
 - modifying QLogic HBA settings 104
 - technical support Web site xvii
 - zoning 35, 75, 95, 132, 165
- system requirements
 - AIX 25
 - HP-UX 65
 - Solaris 85
 - SUSE Linux 124

T

- tasks by documentation title xiv, 159
- trademarks 170

U

- universal transport mechanism
 - See utm device
- upgrading
 - firmware
 - AIX 32
 - HP-UX 72
 - Linux on POWER 127

- upgrading *(continued)*
 - Solaris 92
- NVSRAM
 - AIX 32
 - HP-UX 73
 - Linux on POWER 128
 - Solaris 93

- user tasks
 - DS4000 Storage Manager xiv
 - DS4100 SATA Storage Server 162
 - DS4300 Fibre Channel Storage Server 161
 - DS4400 Fibre Channel Storage Server 160
 - DS4500 Fibre Channel Storage Server 159
 - related documents 163
- utility software
 - package info 7
- utm device 38

V

- variable settings for JNI adapter cards 97
- verify server operation w/ LEDs xii
- VERITAS
 - Dynamic Multipathing (DMP) 7, 107, 108, 110, 113
 - Enterprise Administrator 117
 - File System 107, 113
 - Volume Manager 7, 107, 110, 113
- VolumeCopy 20

W

- Web sites
 - AIX fix delivery center xvi, 37
 - downloading AIX file sets 37
 - DS4000 interoperability matrix xvi
 - DS4000 storage servers xvi
 - DS4000 technical support xvii, 24, 64, 84, 122
 - IBM publications center xvii
 - IBM TotalStorage products xvi
 - Linux on POWER support xvii
 - Linux on pSeries support xvii
 - list xvi
 - premium feature activation xvii, 21
 - readme files xvi, 24, 64, 84, 122
 - SAN support xvii
 - Solaris failover driver info 107
 - Solaris HBA driver info 96, 103
 - switch support xvii

Z

- zoning switches 35, 75, 95, 132, 165

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

IBM TotalStorage DS4000 Storage Manager 9
Installation and Support Guide for AIX, HP-UX, Solaris, and Linux on POWER

Publication No. GC26-7705-00

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

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

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

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

Please tell us how we can improve this book:

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

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

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department GZW
9000 South Rita Road
Tucson, Arizona
USA 85744-0001



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

GC26-7705-00



Spine information:



IBM TotalStorage DS4000 Storage
Manager 9

IBM TotalStorage DS4000 Storage Manager
Installation and Support Guide for AIX, HP-UX,
Solaris, and Linux on POWER