IBM TotalStorage DS4000 Storage Manager Version 9

IBM

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

IBM TotalStorage DS4000 Storage Manager Version 9

IBM

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 227.

Third Edition (June 2005)

Contents

Figures

Tables

About this document

This document provides information about how to set up, install, configure, and work with the IBM® TotalStorage® DS4000 Storage Manager Version 9.1x 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
- Configure the host subsystem
- 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 subsystem. See Appendix E, "Additional DS4000 documentation," on page 215 for a list of these documents.

2. The controller firmware (version 06.1x.xx.xx or later) and NVSRAM that are included with IBM DS4000 Storage Manager Version 9.1x support only the following DS4000 storage subsystems:

   - **DS4100** standard (base) option; Machine Type 1724; Model 100
   - **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
   - **DS4800**; Machine Type 1815; Models 82A/H, 84A/H

   Note: DS4800 storage subsystems are supported only with controller firmware version 06.14.xx.xx.

3. You cannot use controller firmware version 6.1x.xx.xx with the following DS4000 storage subsystems:

   - DS4100 — SCU models
   - DS4300 — SCU models
   - FAStT200
   - FAStT500

   To find the firmware versions that you can use with these storage subsystems, see Table 3 on page 1. Please contact your IBM reseller or representative regarding any future availability of 6.1x.xx.xx controller firmware for these storage subsystems.

FAStT product renaming

IBM has renamed some FAStT family products. Table 1 on page xiv identifies each DS4000 product name with its corresponding previous 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

Previous FAStT Product Name	Current DS4000 Product Name
IBM TotalStorage FAStT Storage Server	IBM TotalStorage DS4000
FAStT	DS4000
FAStT Family	DS4000 Mid-range Disk System
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 xv 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.
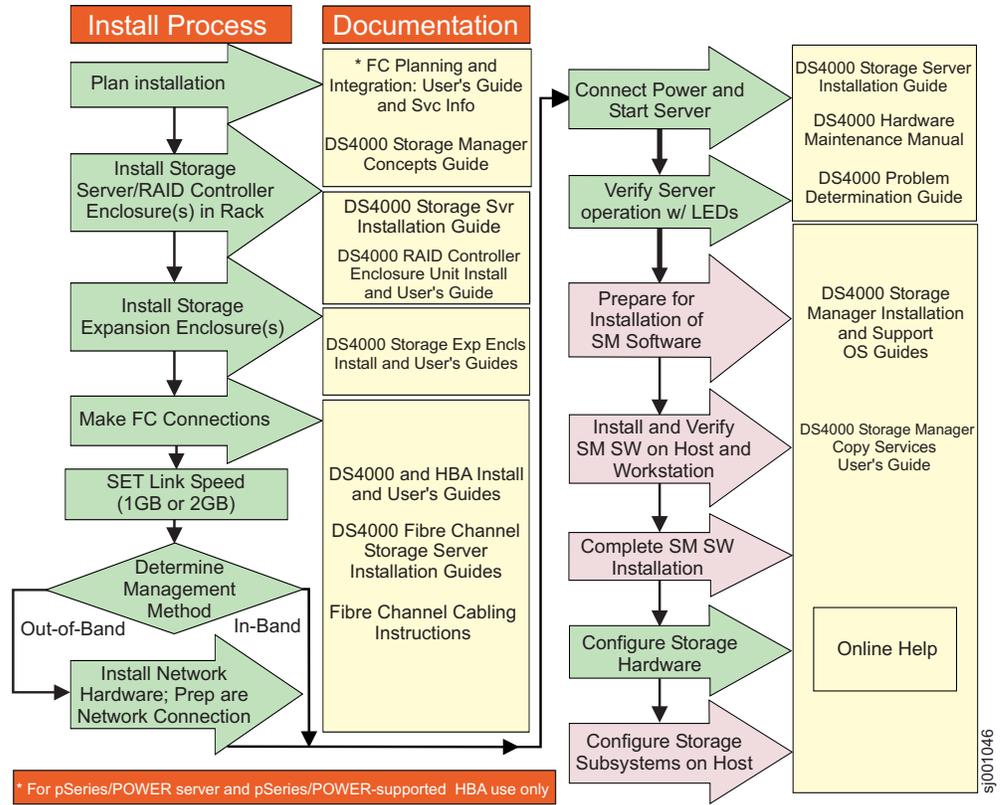
Figure 1. Installation process flow by current publications

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 documentation for these and other DS4000 products at the following Web site:

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

Table 2. TotalStorage DS4000 Storage Manager Version 9.1 titles by user tasks

Title	User tasks					
	Planning	**Hardware installation**	**Software installation**	**Configuration**	**Operation and administration**	**Diagnosis and maintenance**
IBM TotalStorage DS4000 Storage Manager Version 9 Installation and Support Guide for Windows 2000/Server 2003, NetWare, ESX Server, and Linux, GC26-7706	✔		✔	✔		

Table 2. TotalStorage DS4000 Storage Manager Version 9.1 titles by user tasks (continued)

Title	User tasks					
	Planning	Hardware installation	Software installation	Configuration	Operation and administration	Diagnosis and maintenance
IBM TotalStorage DS4000 Storage Manager Version 9 Installation and Support Guide for AIX, UNIX, Solaris and Linux on Power, GC26-7705	✔		✔	✔		
IBM TotalStorage DS4000 Storage Manager Version 9 Copy Services User's Guide, GC26-7701	✔		✔	✔	✔	
IBM TotalStorage DS4000 Storage Manager Version 9 Concepts Guide, GC26-7661	✔	✔	✔	✔	✔	✔

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 27 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 77 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 99 provides information about how to install and configure DS4000 Storage Manager with a Solaris host operating system.

Chapter 5, "Installing and configuring storage management software on POWER-based Linux host systems," on page 133 provides information about how to configure a DS4000 subsystem with a SUSE Linux Enterprise Server 9 (SLES 9) or Red Hat Enterprise Linux 3 (RHEL 3) operating system on a POWER or pSeries-based host.

Chapter 6, "Completing the software configuration," on page 179 provides step-by-step instructions for post-installation tasks, such as how to create storage arrays, logical drives, and storage partitions.

Chapter 7, "Using DS4000 Storage Manager with high-availability cluster services," on page 187 provides information about high-availability clustering system options.

Appendix A, "MC/Service Guard configuration details," on page 199 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 201 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 205 provides the procedures for verifying and downloading Environmental Service Module (ESM) firmware.

Appendix F, "Connecting HBAs in a FC switch environment," on page 223 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 subsystems, 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 subsystems and DS4000 Storage Manager, including documentation and the most recent software, firmware, and NVSRAM downloads, can be found at the following Web sites.

DS4000 Midrange Disk Systems
Find the latest information about IBM TotalStorage disk storage systems, including all of the DS4000 storage subsystems:

www-1.ibm.com/servers/storage/disk/ds4000/

IBM TotalStorage products
Find information about all IBM TotalStorage products:

www.storage.ibm.com/

Support for IBM TotalStorage disk storage systems
Find links to support pages for all IBM TotalStorage disk storage systems, including DS4000 storage subsystems and expansion units:

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

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/disk/ds4000/interop-matrix.html

DS4000 Storage Manager readme files
Find the latest Storage Manager readme files for each host operating system at the IBM TotalStorage support page:

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

Click the link for your DS4000 storage subsystem. When the storage subsystem page opens, click the **Download** tab. Follow the links to the Storage Manager readmes.

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.

@server pSeries and AIX information center
Find everything you need to know about using AIX with pSeries and POWER servers:

publib.boulder.ibm.com/infocenter/pseries/index.jsp?

Linux on pSeries support
Find information about using Linux on pSeries servers:

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

Linux on POWER resource center
Find information about using Linux on POWER servers:

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

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

Premium feature activation
Enable a premium feature on a DS4000 storage subsystem by using the online tool:

www-912.ibm.com/PremiumFeatures/jsp/keyInput.jsp

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 the new features that are incorporated into all releases of controller firmware version 06.1x.xx.xx, and manages DS4000 storage subsystems with all generally-released controller firmware versions (05.xx.xx.xx to 6.1x.xx.xx).

Notes:

1. DS4800 Fibre Channel storage subsystems require controller firmware 06.14.xx.xx.
2. With controller firmware 06.12.xx.xx, Storage Manager 9.1x now supports DS4100 standard (base) SATA storage subsystems and EXP100 SATA storage expansion enclosures.

   DS4100 single-controller (SCU) SATA storage subsystems require controller firmware 05.42.xx.xx.

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

Storage Manager readme files
> Read these first.
>
> For the most recent Storage Manager readme files for your operating system, see the following Web site:
>
> www-1.ibm.com/servers/storage/support/disk/
>
> Click the link for your storage subsystem.
>
> When the page opens, click the **Download** tab, then click the link for Storage Manager.
>
> **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.

All of the controller firmware versions that are listed in Table 3 are available free-of-charge. However, you must purchase a DS4100 single-controller (SCU) storage subsystem to obtain controller firmware 05.42.xx.xx, and you must purchase a DS4800 to obtain controller firmware 06.14.xx.xx.

See Table 3 for information on latest supported controller firmware code levels available for all DS4000 storage subsystem models. See "What's new" on page 5 for new features that are associated with the 06.1x.xx.xx controller firmware.

Table 3. Latest firmware versions by storage subsystem model

Storage subsystem model	Machine type and model	Current firmware version
IBM® TotalStorage® DS4800 Storage Subsystem	1815 - 82A/H and 84A/H	06.14.xx.xx

Table 3. Latest firmware versions by storage subsystem model (continued)

Storage subsystem model	Machine type and model	Current firmware version
IBM TotalStorage DS4500 Storage Subsystem	1742 - 90X/U	06.12.xx.xx
IBM TotalStorage DS4400 Storage Subsystem	1742 - 1RU/RX	06.12.xx.xx
IBM TotalStorage DS4300 Base and Turbo Storage Subsystem	1722 - 60U/X	06.12.xx.xx
IBM TotalStorage DS4300 - Storage Subsystem single controller unit (SCU)	1722 - 60U/X and 6LU/LX	05.34.xx.xx
IBM TotalStorage DS4100 Storage Subsystem	1724 - 100	06.12.xx.xx
IBM TotalStorage DS4100 Storage Subsystem - single controller unit (SCU)	1724 - 100	05.42.xx.xx
IBM FAStT200 Storage Subsystem	3542 - 1RU/RX	05.30.xx.xx
IBM FAStT200 High Availability Storage Subsystem	3542 - 2RU/RX	05.30.xx.xx
Netfinity® FAStT500 RAID Controller Enclosure Unit*	3552 - 1RU/RX	05.30.xx.xx
Netfinity Fibre Channel RAID Controller Unit*	3526 - 1RU/RX	04.01.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 9 (SLES 9) on POWER
- Red Hat Enterprise Linux (RHEL 3) on POWER

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

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 versions of the DS4000 Storage Manager host software, DS4000 storage server controller firmware, DS4000 drive expansion enclosure ESM firmware, and drive 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-1.ibm.com/servers/storage/support/disk/index.html

In the **Additional Support** section of the Web page, 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**. A pull-down menu displays.
3. In the pull-down menu, select **Storage**. Another pull-down menu displays.
4. In the new pull-down menu, and in the subsequent pull-down menus that display, select the following topics:
    - **Computer Storage**
    - **Disk Storage Systems**
    - **TotalStorage DS4000 Midrange Disk Systems & FAStT Stor Srvrs**

    Note: During this process a check list displays. Do not check any of the items in the check list until you complete the selections in the pull-down menus.
5. When you finish selecting the menu topics, 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, then click **Add products**. The **My support** page opens again.
6. On the **My support** page, click the **Edit profile** tab, then click **Subscribe to email**. A pull-down menu displays.
7. In the pull-down menu, select **Storage**. A check list displays.
8. Place a check in each of the following boxes:
    a. **Please send these documents by weekly email**
    b. **Downloads and drivers**
    c. **Flashes**

    and any other topics that you may be interested in, and then click **Update**.
9. Click **Sign out** to log out of My Support.

Overview of DS4000 Storage Manager

IBM TotalStorage DS4000 Storage Manager 9.1x consists of a set of client and host tools that enable you to manage the following IBM DS4000 storage subsystems from a storage management station:

- DS4400 Fibre Channel storage subsystems
- DS4500 Fibre Channel storage subsystems
- DS4300 Standard (Base), Turbo, and Single-Controller (SCU) Fibre Channel storage subsystems
- DS4100 Standard (Base) and Single-Controller (SCU) SATA storage subsystems
- DS4800 Fibre Channel storage subsystems

Notes:

1. DS4800 storage subsystems are supported by Storage Manager 9.14 only (with controller firmware version 6.14.xx.xx).

2. DS4300 and DS4100 Single-Controller (SCU) storage subsystems are not supported with DS4000 controller firmware version 6.1x.xx.xx, or with AIX host operating systems.

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

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

Click the link for your storage subsystem. When the page opens, click the Download tab. Follow the links to the Storage Manager download files.

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 9.1x.

- "FAStT product renaming"
- "New features with firmware 6.14.*xx.xx*"
- "New features with firmware 6.12.*xx.xx* and higher"

FAStT product renaming

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

New features with firmware 6.14.*xx.xx*

DS4000 Storage Manager 9.14 with controller firmware 6.14.xx.xx supports the new DS4800 storage subsystem. If you are using a DS4800 storage subsystem, you must use controller firmware 6.14.xx.xx. Otherwise, use the latest version of controller firmware that is supported for your storage subsystem, as stated in Table 3 on page 1.

Notes:

1. Controller firmware 6.14.xx.xx is supported *only* on the DS4800 storage subsystem.
2. The latest controller firmware versions might be 06.12.xx.xx, 05.42.xx.xx, 05.34.xx.xx, or 5.30.xx.xx, depending on your storage subsystem model.

DS4000 Storage Manager 9.14 supports all of the features listed in "New features with firmware 6.12.*xx.xx* and higher."

New features with firmware 6.12.*xx.xx* and higher

DS4000 Storage Manager version 9.14, with controller firmware 6.12.*xx.xx* and higher, supports the following new features:

New DS4000 Storage Manager installation option
DS4000 Storage Manager 9.14 features an installation wizard that enables you to automatically install Storage Manager software packages on your host server.

Note: Using the DS4000 Storage Manager installation wizard requires a system with a graphics card installed. You still have the option of installing the stand-alone host software packages manually. The packages are included with the installation CD.

DS4000 FC/SATA Intermix update: Premium Feature Key
DS4000 Storage Manager 9.14 supports enabling of the DS4000 FC/SATA Intermix premium feature using a Premium Feature Key.

Important: If you are currently using controller firmware version 6.10.xx.xx with the DS4000 FC/SATA Intermix premium feature enabled, and you want to upgrade to controller firmware version 6.12.xx.xx, see "FC/SATA Intermix premium feature" on page 24 for important information before you install 6.12.xx.xx.

For more information about using the Intermix premium feature, including configuration and set-up requirements, see the *IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview* (GC26-7713).

DS4000 FC/SATA Intermix supported on DS4300 standard (base) subsystems
DS4000 FC/SATA Intermix is supported on DS4300 Standard (Base)
subsystems with controller firmware version 6.12.xx.xx or later.

Support for DS4100 standard (base) subsystems
Storage Manager 9.14, with controller firmware 06.12.xx.xx or later,
supports DS4100 standard (base) subsystems.

> **Note:** The VolumeCopy and FC/SATA Intermix premium features are not
> supported at this time with the DS4100 Standard (Base) storage
> subsystem. Also, the DS4100 Standard (Base) storage subsystem is
> not supported on AIX host operating systems.

Enhanced Remote Mirroring supported on DS4100 standard (base) subsystems
Enhanced Remote Mirroring is now supported on DS4100 standard (base)
subsystems with controller firmware version 6.12.xx.xx or later.

Enhanced Remote Mirroring
In addition to Metro Mirroring, IBM DS4000 Storage Manager version 9.1x
and higher, with controller firmware level 6.10.11.xx or higher, also supports
Global Copy and Global Mirroring Remote Mirror options. Please see the
*IBM TotalStorage DS4000 Storage Manager Version 9 Copy Services
User's Guide* for more information.

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

Parallel hard drive firmware download
You can now download drive firmware packages 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.

> **Notes:**
> 1. Drive firmware download is an offline management event. You must
>   schedule downtime for the download because I/O to the storage
>   subsystem is not allowed during the drive firmware download process.
> 2. Parallel hard drive firmware download is not the same thing as
>   concurrent download.

Staged controller firmware download
You can now download the DS4000 controller firmware and NVSRAM to
DS4300 Turbo and DS4500 storage subsystems for later activation.
Depending on your firmware version, DS4000 storage subsystem 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 subsystems.

Subsystem Management Window usability enhancements
Storage Manager 9.1x features the following usability enhancements:

- Storage Partitioning wizard, which helps you easily create storage partitions
- Task Assistant, which helps guide you through common enterprise and subsystem management tasks
- Ability to extract SMART data for SATA drives
- One-click collection of support data, drive diagnostic data, drive channel state management, controller 'service mode,' and the ability to save host topology information
- Improved media error handling for better reporting of unreadable sectors in the DS4000 storage subsystem event log, and persistent reporting of unreadable sectors

In addition, troubleshooting, recovery and maintenance tools are now under the Advanced heading in the Storage Manager 9.1x Subsystem Management window. The following submenus are available:

- Maintenance
- Troubleshooting
- Recovery

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.

Support for DS4300 standard (base) storage subsystems
Controller firmware 06.12.xx.xx and higher supports DS4300 Standard (Base) storage subsystems.

> **Note:** The VolumeCopy, FC/SATA Intermix and Enhanced Remote Mirroring premium features are not supported at this time with the DS4300 Standard (Base) storage subsystem. Also, the DS4300 Standard (Base) storage subsystem is not supported on AIX host operating systems.

Support for DS4000 EXP710 storage expansion enclosures
Controller firmware 06.1x.xx.xx and higher supports supports DS4000 EXP710 storage expansion enclosures.

Increased support for DS4000 EXP100 SATA storage expansion enclosures
DS4000 EXP100 SATA storage expansion enclosures are now supported on DS4400 storage subsystems

Also, the DS4100 storage subsystem now supports up to seven EXP100 SATA storage expansion enclosures.

New host operating system information
The following host features are now available with DS4000 Storage Manager:

- **Support for DS4000 Storage Manager on SUSE Linux 9**

  You can install and run DS4000 Storage Manager 9.14 host software on a SUSE Linux 9 (SLES 9) host, with Service Pack 1.

Note: If you are using a SLES 9 host with a Storage Manager version earlier than 9.14, you must use a different operating system for your management station.

- **Support for RDAC on SUSE Linux 9**

  POWER-based SUSE Linux 9 (SLES 9) hosts now support the RDAC multi-path failover driver.

- **Support for RDAC on RHEL 3**

  POWER-based Red Hat Enterprise Linux 3 (RHEL 3) hosts now support the RDAC multi-path failover driver.

- **Support for DS4000 Storage Manager on Red Hat Enterprise Linux 3**

  You can install and run DS4000 Storage Manager client software on a POWER-based Red Hat Enterprise Linux 3 (RHEL 3) host, and use that host as your storage management station.

DS4000 limitations

Please note the following limitations:

- DS4000 controller firmware version 06.14.xx.xx supports only DS4800 storage subsystems.
- DS4100 and DS4300 single-controller (SCU) storage subsystems do not support DS4000 controller firmware 6.1x.xx.xx.
- DS4800 storage subsystems do not support DS4000 EXP700 Fibre Channel storage expansion enclosures. You must upgrade to EXP710 storage expansion enclosures.
- Controller firmware version 06.1x.xx.xx and corresponding NVSRAM are not supported on the following storage subsystems:
  - FAStT200 (3542, all models)
  - FAStT500 (3552, all models)

  Therefore, the Storage Manager features that are specific to version 06.1x.xx.xx firmware are unavailable for these FAStT subsystems. However, IBM recommends that you do use Storage Manager 9.1x software to manage these FAStT subsystems, using controller firmware versions 5.30.xx.xx. For more information about which firmware to use with your configuration, see Table 3 on page 1.
- DS4300 Turbo storage subsystems do not support DS4000 EXP500 Fibre Channel storage expansion enclosures.
- DS4100 and DS4300 standard (base) storage subsystems do not support the FC/SATA Intermix premium feature.

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

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 subsystems 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 subsystems
- 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 to manage storage subsystems through the host Fibre Channel connection.

Notes:

1. You must have a failover driver installed in order to use SMagent.
2. 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

SMutil software package

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

Install SMutil on all HP-UX, Solaris, and Linux host servers that are attached to a storage subsystem. (The host servers are attached to the storage subsystem through the Fibre Channel.)

Failover drivers

AIX, Linux on POWER, and Solaris host systems require one of the following failover drivers for Fibre Channel path redundancy:

AIX, RHEL 3, SLES 9
> 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.

Notes:

1. 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 and follow the instructions in that section.
2. HP-UX hosts do not require a nonstandard failover driver for multipath configurations. With the standard HP-UX driver, you can enable multipath I/O by using physical volume (PV) links.

Software installation sequence

Install the DS4000 Storage Manager software packages in the sequences shown in

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

Step	AIX	HP-UX	Solaris	Linux on POWER
1	SMruntime	SMruntime	SMruntime	SMruntime
2	SMclient	SMclient	SMclient	SMclient
3	RDAC	SMagent	SMutil	SMutil
4	SMagent	SMutil	RDAC or DMP	RDAC
5			SMagent	SMagent
Location of procedures	"Installing DS4000 Storage Manager" on page 31	"Installing DS4000 Storage Manager" on page 80	"Installing DS4000 Storage Manager" on page 102	"Installing DS4000 Storage Manager and utilities" on page 143

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 12 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.

Storage subsystems

Controller

Controller

Running the
host-agent software

Host computer

Storage subsystems

Controller

Controller

Fibre Channel
I/O path

Network

Note: The host can also act as a
management station.

Management station
(one or more)

SJ000707
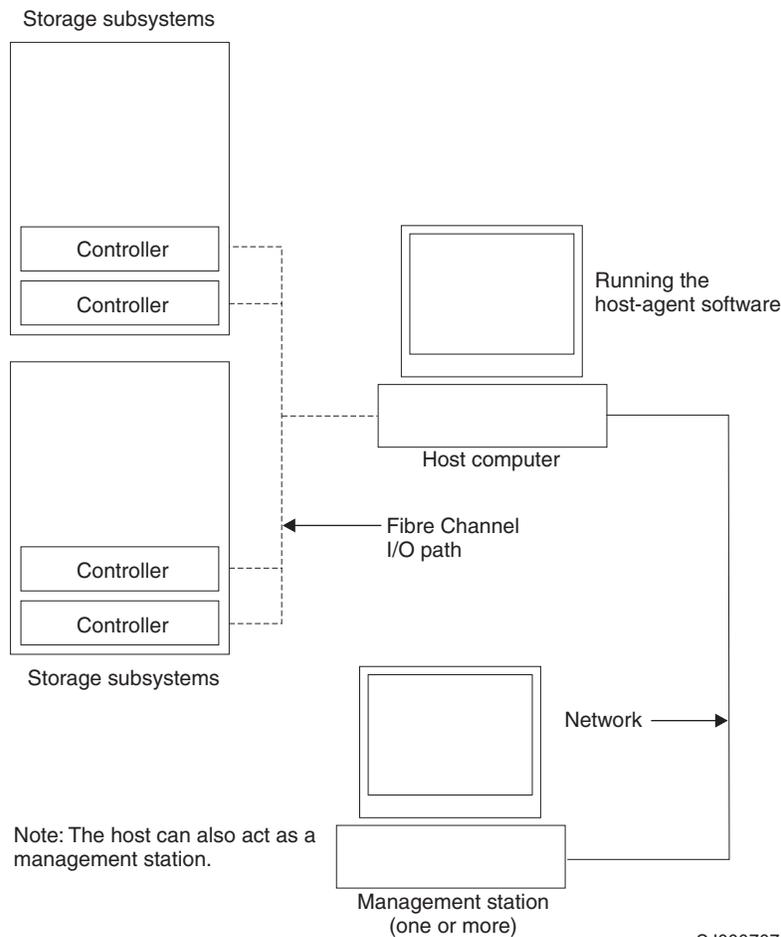
Figure 2. In-band managed storage subsystems

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 (as described in the installation chapters for each operating system).
- 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 14 shows the out-of-band management method.
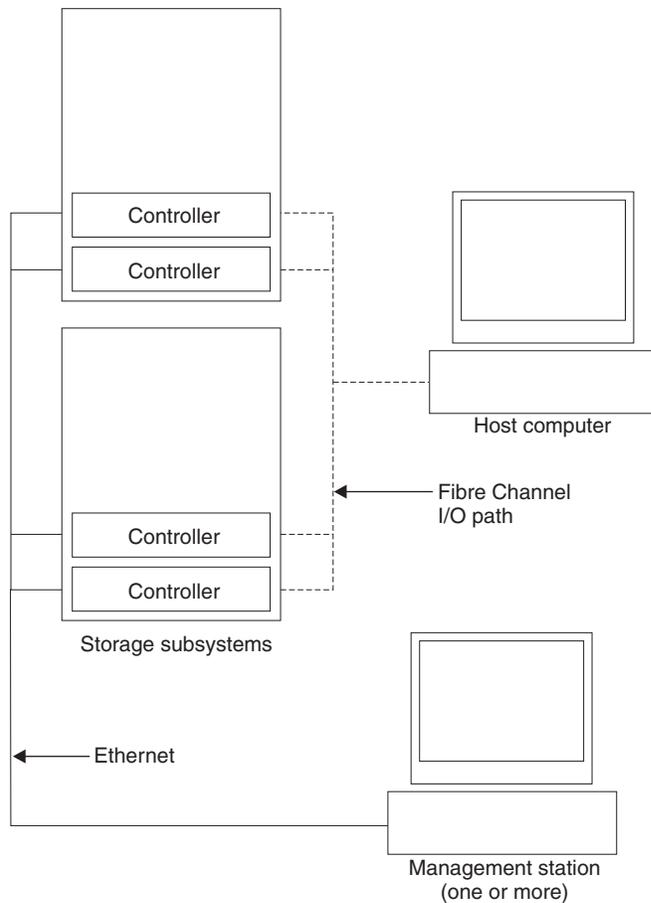
Figure 3. Out-of-band managed storage subsystems

If you use the out-of-band management method, see "Setting up IP addresses for DS4000 storage controllers" on page 20.

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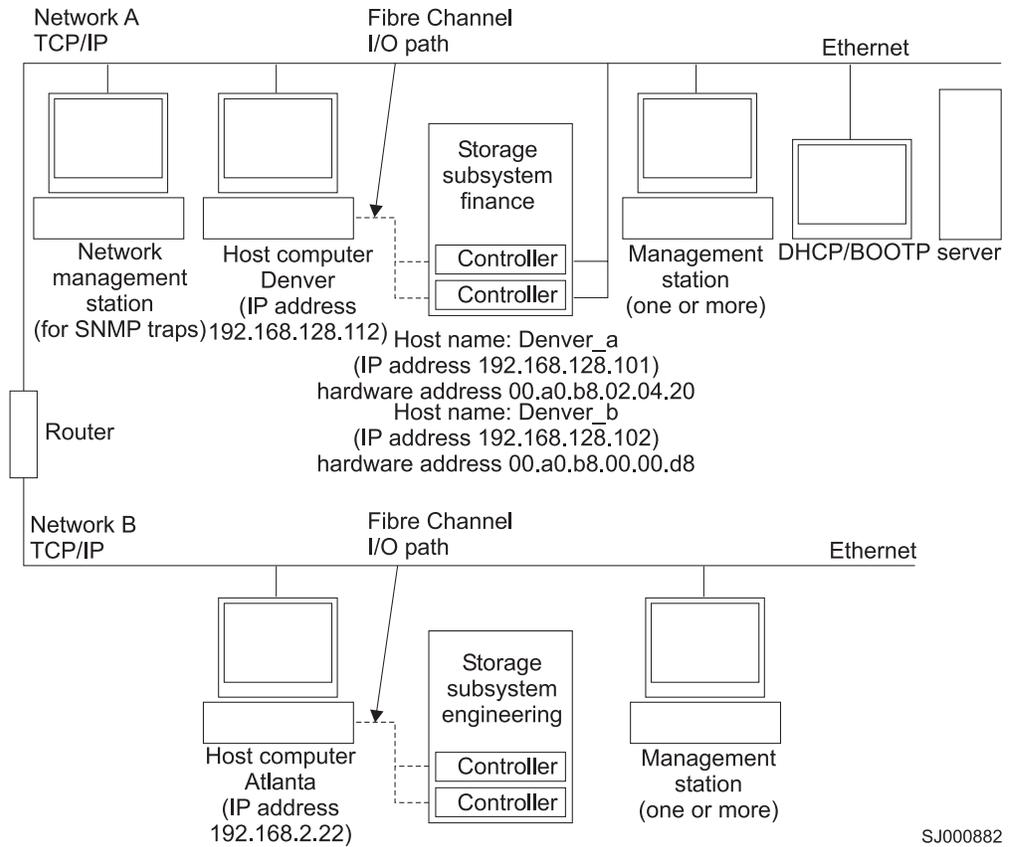
Figure 4. Sample network using both out-of-band and in-band managed storage subsystems

Installation process diagram

Figure 5 on page 17, Figure 6 on page 18, and Figure 7 on page 19 provide an overview of the installation process for out-of-band management.

```
                    ┌──────────────────┐
                    │  Set up DS4000   │
                    │    Hardware      │
                    └──────────────────┘
                             │
                    ┌──────────────────┐
                    │ Determine Fibre  │
                    │ Channel connection│
                    │     method       │
                    └──────────────────┘
                             │
                          ◇ Direct-
             Yes          Attach?           No
```

Set up DS4000 Hardware

Determine Fibre Channel connection method

Direct-Attach?

Yes — Make FC connection directly to controller

No — Verify and update switch firmware

Create zones on FC switch or switches

Make FC connection using FC switch

Configure network settings on DS4000 controllers

Install SM Client

Windows or UNIX?

Windows — Install SM Client on Windows mgmt station

UNIX — Install SM Runtime on UNIX host

Install SM Client on UNIX host

(Continued on Flowchart 2)   SJ002000

Figure 5. Installation process (part 1)

(Continued from Flowchart 1)

```
              ┌─────────────────────┐
              │ Verify and update   │
              │ controller firmware │
              │ and NVSRAM          │
              └──────────┬──────────┘
                         │
                    ╱────┴────╲
       Yes        ╱  Install   ╲        No
  ┌──────────────│   Premium    │──────────────┐
  │              │   Feature    │              │
  │              ╲    Keys     ╱               │
  │               ╲────┬────╱                  │
┌─┴──────────────┐     │                       │
│ Enable feature │     │                       │
└─┬──────────────┘     │                       │
  │                    │                       │
  └────────────────────┴───────────────────────┘
                       │
              ┌────────┴────────┐
              │ Configure arrays │
              │ and create logical│
              │ drives          │
              └────────┬────────┘
                       │
              ┌────────┴────────┐
              │   Host Setup    │
              └────────┬────────┘
```

Solaris	AIX	HP-UX
Check patch levels on Solaris host	Check HBA firmware levels	Check patch levels on HP-UX host
Prepare host for host software	Check PTF/APAR levels on AIX host	Prepare host for host software
	Check levels for HBA and RDAC device drivers	

Solaris branch:

- **SM Runtime installed?** Yes / No
  - No → Install SM Runtime on Solaris host
- Install SM Agent on Solaris host
- Install SM Util on Solaris host
- Install HBA driver
- Set HBA bindings

HP-UX branch:

- **SM Runtime installed?** Yes / No
  - No → Install SM Runtime on HP-UX host
- Install SM Agent on HP-UX host
- Install SM Util on HP-UX host
- Reboot HP-UX host

(Continued on Flowchart 3)

SJ002001

Figure 6. Installation process (part 2)

(Continued from Flowchart 2)

Direct Attach?

Yes

No

Note: Direct Attach does not require bindings to be set

Modify HBA Bindings

Install RDAC (Failover driver) on Solaris Host

Partition Storage?

Yes

No

Define Host Groups

Define Host

Define Both Host-Ports

Map Logical Drives to Partition

End
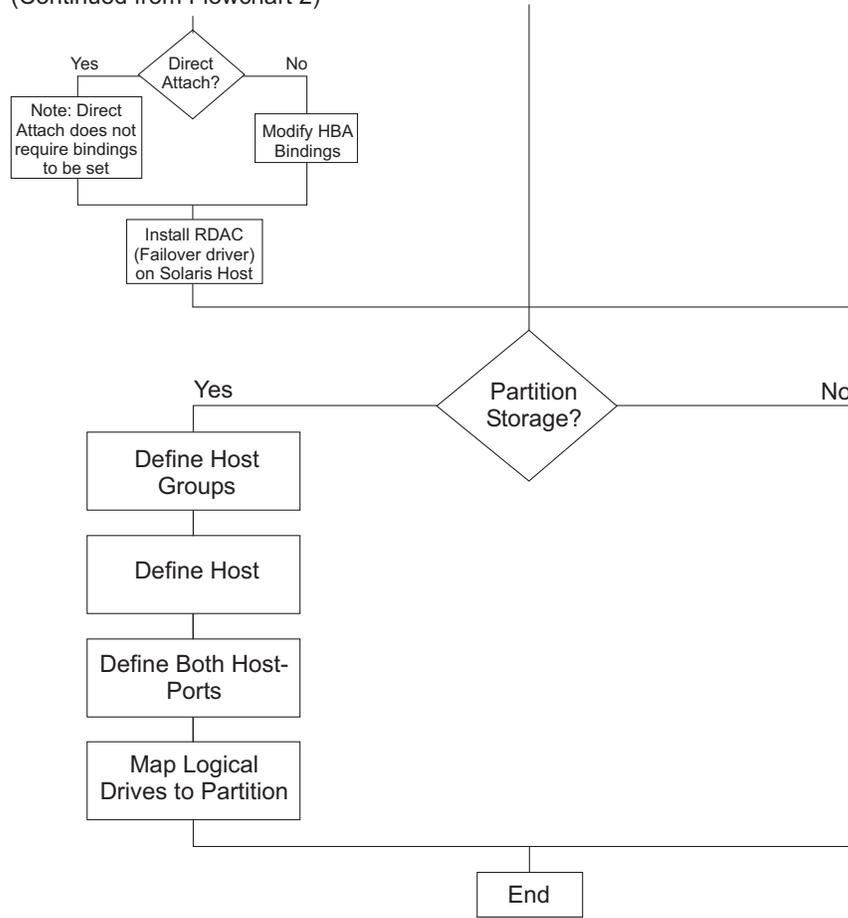
SJ002002

Figure 7. Installation process (part 3)

Setting up IP addresses for DS4000 storage controllers

Complete the following procedures after you install SMruntime and SMclient, as described in the installation section for your host operating system:

1. "Setting up the DHCP/BOOTP server and network"
2. "Setting up the DHCP/BOOTP server and network"

To complete these procedures, you must set up a DHCP or BOOTP server and network with the following components:

- A DHCP or 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: You can avoid DHCP/BOOTP server and network tasks by assigning static IP addresses to the controller. If you do not wish to assign static TCP/IP addresses with the DS4000 Storage Manager, using the DS4000 default TCP/IP addresses, see the IBM support Web site at:

www.ibm.com/pc/support/

Refer to Retain Tip # H171389 Unable To Setup Networking Without DHCP/BOOTP.

Setting up the DHCP/BOOTP server and network

Complete the following steps to set up the DHCP/BOOTP server and network:

1. Get the MAC address from each controller blade.
2. Complete whichever of the following steps is appropriate for your server:
   - On a DHCP server, create a DHCP record for each of the MAC addresses. Set the lease duration to the longest time possible.
   - On a BOOTP server, edit the bootptab file to add in the entries that associate the MAC address tab with the TCP/IP address.
3. Connect the DS4000 storage subsystem Ethernet ports to the network.
4. Boot the DS4000 storage subsystem.

When you are finished, complete the steps in the next section, "Steps for assigning static TCP/IP addresses to the DS4000 controllers."

Steps for assigning static TCP/IP addresses to the DS4000 controllers

Complete the following steps to assign static TCP/IP addresses to the DS4000 storage subsystem controllers, using default TCP/IP addresses that are assigned to the DS4000 storage subsystem controllers during manufacturing:

1. Make a direct management connection to the DS4000 storage subsystem, using the default TCP/IP addresses:
   - **Controller A:** 192.168.128.101
   - **Controller B:** 192.168.128.102
2. Start SMclient. The Enterprise Management window opens.
3. In the Enterprise Management window, click on the name of the default storage subsystem. The Subsystem Management window opens.

4. In the Subsystem Management window, right-click the controller icon and select **Change —> Network Configuration** in the pull-down menu. The Change Network Configuration window opens.

5. In the Change Network Configuration window, click on the Controller A and Controller B tabs and type the new TCP/IP addresses in their appropriate fields. Click **OK**.

6. Close the Subsystem Management window, wait five minutes, then delete the default DS4000 storage subsystem entry in the Enterprise Management window.

7. Add a new storage subsystem entry in the Enterprise Management window, using the new TCP/IP address.

Note: To manage storage subsystems through a firewall, configure the firewall to open port 2463 to TCP data.

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 subsystem, 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
HP-UX 11.i	128
Solaris	256 with DMP 32 with RDAC
SUSE Linux Enterprise Server 9	256
Red Hat Enterprise Linux 3	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 controller 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 controller 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	DS4000 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 6.1x.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 9.1x:
> - FlashCopy®
> - VolumeCopy
> - Enhanced Remote Mirror Option

Storage Partitioning premium feature
> Storage Partitioning is standard on all DS4000 storage subsystems that are supported by DS4000 controller firmware version 6.1x.xx.xx.

"FC/SATA Intermix premium feature" on page 24
> The IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix

premium feature supports the concurrent attachment of Fibre Channel and SATA storage expansion enclosures to a single DS4000 controller configuration.

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 9 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 DS4000 controller firmware version 5.42.xx.xx, or with the DS4100 storage subsystem.

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 DS4000 controller firmware versions earlier than 6.1x.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 DS4000 controller firmware versions 5.41.xx.xx or 5.42.xx.xx, or with the DS4100 storage subsystem.

Storage Partitioning premium feature

The Storage Partitioning feature enables you to associate a set of logical drives on a DS4000 storage subsystem 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 subsystem.

Table 7. Using Storage Partitioning with DS4000 storage subsystems

DS4000 Storage subsystem	Storage Partitioning enabled by default	Maximum number of storage partitions
DS4800 Models 1815-82H and 84H	Yes (8 partitions standard)	64
DS4800 Models 1815-82A and 84A	Yes (Standard number depends upon configuration)	64
DS4100 Standard (Base) and Single-Controller (SCU)	No	16
DS4500	Yes (16 partitions standard)	64
DS4400	Yes (64 partitions standard)	64
DS4300 with Turbo option	Yes (8 partitions standard)	64
DS4300 Standard (Base) and Single-Controller (SCU)	No	16

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 subsystem 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 subsystem, access to all logical drives might be satisfactory.

Note: For more information about fabric switches, see Appendix F, "Connecting HBAs in a FC switch environment," on page 223.

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

FC/SATA Intermix premium feature

The IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix premium feature supports the concurrent attachment of Fibre Channel and SATA storage expansion enclosures to a single DS4000 controller configuration.

With this premium feature, IBM TotalStorage DS4500, DS4400, and DS4300 Turbo storage subsystems can support the attachment of DS4000 EXP700 and EXP710 Fibre Channel storage expansion enclosures, and DS4000 EXP100 Serial ATA storage expansion enclosures, concurrently.

In this way, the Intermix premium feature enables you to create and manage distinct arrays that are built from either Fibre Channel disks or SATA disks, and allocate logical drives to the appropriate applications using a single DS4000 storage subsystem.

Please see your IBM technical service representative for information regarding future DS4000 storage subsystem support for the FC/SATA Intermix premium feature.

For important information about using the Intermix premium feature, including configuration and set-up requirements, see the *IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview* (GC26-7713).

Enabling the DS4000 FC/SATA Intermix premium feature

Storage Manager 9.1x supports the DS4000 FC/SATA Intermix premium feature with all 06.1x.xx.xx controller firmware versions. However, depending on which level of the controller firmware is installed on your storage subsystem, you need to follow different procedures to enable the premium feature.

DS4000 controller firmware version 6.10.xx.xx
> On storage subsystems using DS4000 controller firmware version 6.10.xx.xx, the DS4000 FC/SATA Intermix premium feature is enabled using an NVSRAM file.
>
> Download the appropriate NVSRAM version, as described in the *IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview*, Second Edition (GC26-7713-01), to enable the DS4000 FC/SATA Intermix premium feature.

DS4000 controller firmware versions 6.12.xx.xx and 6.14.xx.xx
> On storage subsystems using DS4000 controller firmware versions 6.12.xx.xx and 6.14.xx.xx, the DS4000 FC/SATA Intermix premium feature is enabled using a premium feature key.
>
> Install the DS4000 FC/SATA Intermix Premium Feature Key, as described in the *IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview*, Third Edition (GC26-7713-02), to enable the DS4000 FC/SATA Intermix premium feature.

Upgrading DS4000 controller firmware on systems with FC/SATA Intermix already enabled: If you are currently using DS4000 controller firmware 6.10.xx.xx and want to upgrade to 6.12.xx.xx or 6.14.xx.xx, and you already have the FC/SATA Intermix premium feature enabled on your subsystem, you need to complete the steps in this section.

Complete the following steps to upgrade the controller firmware to 6.12.xx.xx or 6.14.xx.xx:

1. Use the instructions in the DS4000 FC/SATA Intermix premium feature activation card to generate the premium feature key file.
2. Upgrade the controller firmware and NVSRAM to 06.12.xx.xx or 6.14.xx.xx version.

> **Note:** After the firmware is upgraded, the recovery guru will display the error "premium feature out-of-compliance." The premium feature that is out-of-compliance is the DS4000 FC/SATA Intermix premium feature. You will not be able to create any arrays or logical drives or modify the LUNs until after you complete step 3.

3. Enable the DS4000 FC/SATA Intermix premium feature using the premium feature key file. This restores the DS4000 subsystem to full functionality.

Enabling your premium features

Complete the following steps to enable a premium feature:

Note: Before you enable the FC/SATA Intermix premium feature, see "FC/SATA Intermix premium feature" on page 24, and the *IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview* (GC26-7713).

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-912.ibm.com/PremiumFeatures/jsp/keyInput.jsp

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 30
- "Installing DS4000 Storage Manager" on page 31
- "Configuring storage subsystems" on page 34
- "Creating direct and SAN-attached configurations" on page 41
- "Installing the RDAC failover driver" on page 42
- "Identifying devices on AIX hosts" on page 44
- "Verifying the installation and configuration" on page 46
- "Viewing and setting Object Data Manager (ODM) attributes" on page 55
- "Additional AIX configuration information" on page 62
- "Replacing hot swap HBAs" on page 66
- "Resolving disk array errors" on page 72
- "Redistributing logical drives in case of failure" on page 75

Hardware, software and host system requirements

This section lists 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 Storage Manager readme file for AIX that is located in the 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 subsystem. When the storage subsystem page opens, click the **Download** tab, and follow the links to the readmes.

DS4000 hardware and firmware requirements

Table 8 lists the DS4000 Fibre Channel storage subsystems that you can use with DS4000 Storage Manager.

Table 8. DS4000 storage subsystems that you can use with AIX systems

Product name	Model	Controller firmware versions
IBM TotalStorage DS4800 Storage Subsystem	1815-82A/H, 1815-84A/H	06.14.xx.xx
IBM TotalStorage DS4100 Storage Subsystem	1724-100	06.12.xx.xx
IBM TotalStorage DS4500 Storage Subsystem	1742-90U	06.12.xx.xx

Product name	Model	Controller firmware versions
IBM TotalStorage DS4400 Storage Subsystem	1742-1RU	06.12.xx.xx
IBM TotalStorage DS4300 Standard and Turbo Storage Subsystem	1722-60U, 1722-60X	06.12.xx.xx

NVSRAM: For the latest NVSRAM versions, see the most recent Storage Manager readme files for AIX at the following Web site:
www-1.ibm.com/servers/storage/support/disk/
Click the link for your storage subsystem. When the storage subsystem page opens, click the **Download** tab, then click the Storage Manager link, then click the **Storage Mgr** tab.

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 Enclosure	IBM TotalStorage DS4000 EXP100 SATA Storage Expansion Enclosure	1710-10U or 10X	9554 **Note:** When upgrading the EXP100 to ESM firmware version 9554, you must use a storage subsystem that has DS4000 controller firmware **05.41.1x.xx** installed . You cannot upgrade to ESM firmware 9554 using a storage subsystem with different firmware installed.
	IBM TotalStorage DS4000 EXP710 Fibre Channel Storage Expansion Enclosure	1740-710	9627
	IBM TotalStorage DS4000 EXP700 Fibre Channel Storage Expansion Enclosure	1740-1RU or 1RX	9326 **Note:** You must upgrade the EXP700 ESM firmware to version 9326 before you upgrade the DS4000 controller firmware to version **06.1x.xx.xx**.
	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 5.2 and 5.3 only)	FC 5716	1.90AX4
Switches	IBM TotalStorage SAN Switch	2109-S08 2109-S16	2.6.2A
	IBM TotalStorage SAN Switch	2109-F16	3.1.2A
	IBM TotalStorage SAN Switch	2109-M12	4.0.2C
	IBM TotalStorage SAN Switch	2109-F32	4.2.0C
	IBM TotalStorage SAN 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
	CNT switch	2042-001	4.3.0
	Cisco switch	2062-D01, 2062-D07, 2062-T07	1.3.4A

Note: Check the latest readme file for the most up-to-date firmware versions.

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 subsystem. When the storage subsystem page opens, click the **Download** tab.

Note: If you have any questions about the latest DS4000 interoperability information, see the DS4000 interoperability matrix at the following Web site:

www-1.ibm.com/servers/storage/disk/ds4000/interop-matrix.html

AIX system requirements

This section lists the minimum hardware and software requirements for an AIX host in a DS4300, DS4400, DS4500, or DS4800 storage subsystem that is managed by DS4000 Storage Manager.

For more information about using AIX with pSeries and POWER servers, see the pSeries and AIX Information Center at the following Web site:

publib.boulder.ibm.com/infocenter/pseries/index.jsp?

Hardware requirements

The following hardware is required at minimum:

Any RS/6000®, RS6000 SP or IBM @server pSeries™ or POWER server that supports FC 6227, FC 6228, FC 6239 or FC 5716 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 system firmware is required.

Software requirements

The following operating system software is required:

- AIX 5.1, 5.2, 5.3, or Virtual I/O Server
- File sets associated with IBM RDAC driver installations that are listed in Table 11 on page 43, Table 12 on page 43, or Table 13 on page 43.

Note: For information about Virtual I/O Server requirements, see the following Web site: techsupport.services.ibm.com/server/virtualization/vios

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 subsystem. You can configure up to two HBAs per partition and up to two partitions per DS4000 storage subsystem. Additional HBAs can be added for additional DS4000 storage subsystems and other SAN devices, up to the limits of your specific subsystem platform.
- Direct-attach configurations are restricted to single-initiator configurations only. You can only make one connection to each minihub on a DS4000 storage subsystem.

  Note: DS4800, DS4300, and DS4100 storage subsystems 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.
- 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.

  Note: DS4800, DS4300, and DS4100 storage subsystems do not support daisy-chain configurations.

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 5.2 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 5.2, 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 subsystem 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 a SATA device.

Partitioning restrictions

- The maximum number of partitions per AIX host, per DS4000 storage subsystem, is two.

- All logical drives that are configured for AIX must be mapped to an AIX host group.

  Note: If you are using the default host group, ensure that the default host type is AIX.
  For more information, see "Storage Partitioning: Defining an AIX host group" on page 38 and "Storage Partitioning premium feature" on page 24.

- 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 dynamic volume expansion (DVE) on AIX 5.1.

- DS4100 and DS4300 *single-controller* storage subsystems are not supported with AIX hosts. (DS4100 and DS4300 dual-controller storage subsystems are supported with AIX hosts.)

- EXP700 storage expansion units are not supported with DS4800 storage subsystems. You must upgrade to EXP710 storage expansion units.

Installing DS4000 Storage Manager

You can install all of the DS4000 Storage Manager software packages automatically by using the Storage Manager installation wizard, or you can install each package manually. This section contains the following procedures:

- "Installing Storage Manager using the installation wizard"
- "Installing Storage Manager packages manually" on page 32 (original method)

Installing Storage Manager using the installation wizard

The DS4000 Storage Manager installation wizard installs the following software packages on the AIX host:

- SMruntime
- SMclient
- SMagent
- SMutil

Requirements:
1. DS4000 Storage Manager 9.12 or 9.14
2. Graphics adapter card

If you are installing Storage Manager 9.10, or if your system does not have a graphics card installed, or if for any reason you do not want to use the wizard to install the software, skip this section and install the stand-alone host software packages using the procedures described in "Installing Storage Manager packages manually." All the packages are included with the installation CD.

Notes:
1. If you use the wizard to install Storage Manager, you might need to increase the size of some partitions to accommodate the installation.
2. Do not install the software packages in the same subsystem by using both the installation wizard and the manual installation method.

Complete the following steps to install the DS4000 Storage Manager software using the installation wizard. Adjust the steps as necessary for your specific installation.
1. Download the following file from the DS4000 Storage Manager CD, or from the DS4000 support Web site, to a directory on your system:

   SMIA-AIX-09.14.xx.xx.bin
2. Open the file on your system. The Storage Manager Installation wizard's Introduction window opens.
3. Follow the instructions in each window of the wizard. When you select the installation type, you can choose one of the following options:
   * Typical (Full) Installation — Installs all Storage Manager software packages
   * Management Station — Installs SMruntime and SMclient
   * Host — Installs SMruntime, SMagent and SMutil
   * Custom — Allows you to select which packages you want to install
4. Restart the system.

Result: The DS4000 Storage Manager software is installed on your system.

Installing Storage Manager packages manually

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 33

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 27 and "Software requirements" on page 30.

- 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 42 to download them from the appropriate Web site. For lists of required file sets, see the following tables in "AIX system requirements" on page 29:

  – **AIX 5.1:** Table 11 on page 43
  – **AIX 5.2:** Table 12 on page 43
  – **AIX 5.3:** Table 13 on page 43

- 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-09.1x.xx.xx.bff SMruntime.aix.rte
```

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

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

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

The verification process returns a table that describes the software installation, including the install package file name, version number, action and action status. If the verification process returns an error, contact your 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-09.1x.xx.xx.bff SMclient.aix.rte
```

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

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

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

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

Installing 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 29.
- SMruntime has been installed.

Note: SMagent is not supported with DS4000 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-9.1x.xx.xx.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:

```
# lslpp -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 26.

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 20 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 Storage Subsystem**.

    b. In the Add Storage Subsystem window, type the IP addresses of each controller in the storage subsystem, and click **Add**. The name of the storage subsystem displays in the Enterprise Management window.

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

    The name of the storage subsystem displays in the Enterprise Management window.

4. Double-click the name of the storage subsystem to open its Subsystem Management window.

Upgrading controller firmware and NVSRAM

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

Overview

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

- "Using concurrent download"
- "Using traditional or staged download"

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 55 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.

    To find out whether you can use concurrent download, see Table 10 on page 37.

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 DS4000 controller firmware version 5.4x.xx.xx to version 6.1x.xx.xx. You can only use staged download if the DS4000 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 subsystems.

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 subsystem 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 subsystem that you plan to update.

Preparing for firmware or NVSRAM download

To prepare for download, review and complete all of the steps in this section before you start the download procedures.

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

To check the controller 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 view the current versions of controller firmware and NVSRAM.
3. Click **Close** to close the Storage Subsystem Profile window.
4. To find whether the versions of controller firmware and NVSRAM that are currently on your storage subsystem are the latest versions, see the Storage Manager readme file for AIX at the following Web site:

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

   Click the link for your storage subsystem. When the storage subsystem page opens, click the **Download** tab.

   You can also find controller firmware versions listed in Table 3 on page 1.
5. If you find that you do not have the latest versions of the firmware and NVSRAM, download the correct versions from this Web site to a directory on your system, then follow the preparation and upgrade procedures that are described in this section.

Save the storage subsystem configuration: Saving the storage subsystem configuration to a script file is necessary in order to restore the configuration in the event of a catastrophic failure.

Attention: Do not use this option if the storage subsystem is undergoing an operation that changes any configuration settings. (If any logical drive listed in the Logical View of the Subsystem Management window has a clock icon, then a configuration change is occurring, and you must wait to save the configuration.)

You can also save the storage subsystem profile, by selecting **Storage Subsystem —> View Profile** in the Subsystem Management window, and clicking **Save As**. Then, in the Save Profile window, select the sections of the profile that you want to save, specify a name and location for the text file, and click **Save**.

Save the storage subsystem configuration by completing the following steps:

1. In the Subsystem Management window, click **Storage Subsystem —> Configuration —> Save**. The Save Configuration window displays.
2. In the Save Configuration window, select the elements of the configuration that you want to save, and click **Yes**. A save dialog displays.
3. In the save dialog, specify a name for the script file and select its location.

   Notes:

   - The .cfg extension is used by default for the script file if you do not specify a file extension.
   - Do not store your DS4000 Storage Subsystem configuration script file in the same location as the logical drives that are defined on your DS4000 Storage Subsystem. If you do, and the logical drives fail during the subsequent migration, you will lose data.

4. Click **Save**. The Configuration Saved window displays.
5. Review the information in the Configuration Saved window, then click **OK** to return to the Subsystem Management window.

Determine whether you can use concurrent download: You can use concurrent download only if you are upgrading from one version of controller firmware to a later release of the same version. For example, if you are upgrading from 06.10.xx.xx to 6.12.xx.xx, you can use concurrent download. However, if you are upgrading from 5.4x.xx.xx to 06.1x.xx.xx, you *cannot* use concurrent download, and you need to stop I/O before upgrading controller firmware.

Table 10. Firmware versions that support concurrent download (AIX)

DS4000 storage subsystem:	Controller firmware version:
DS4300 base and Turbo:	6.10.xx.xx. 6.12.xx.xx
DS4400/DS4500:	6.10.xx.xx, 6.12.xx.xx
DS4800	06.14.xx.xx
DS4100	06.10.xx.xx. 06.12.xx.xx
All other DS4000 storage subsystems:	05.30.xx.xx, 6.10.xx.xx, 6.12.xx.xx

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

1. Download the correct version of the firmware from the Web site that is listed in "Check current versions of controller firmware and NVSRAM" on page 36. Place the file in a designated directory on the host system.

2. Check your system's current firmware and device driver levels to determine whether you need to stop I/O before upgrading the firmware, using the information that is listed in Table 10 on page 37.

   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.

3. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.

4. In the Subsystem Management window, click **Advanced —> Maintenance —> Download —> Controller Firmware**. The Download Firmware window displays.

5. Click **Browse** to select the firmware file that you previously downloaded to a directory on the host.

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**. The Confirm Download window displays.

8. Click **Yes** to start the download.

9. Close the Subsystem Management window and then reopen it to complete the firmware update.

10. Verify that the firmware was successfully installed.

Steps for upgrading the NVSRAM

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

1. Download the correct NVSRAM version from the Web site that is listed in "Check current versions of controller firmware and NVSRAM" on page 36. Place the file in a designated directory on the host system.

2. Stop I/O.

3. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.

4. In the Subsystem Management window, click **Advanced —> Maintenance —> Download —> Controller NVSRAM**. The Download NVSRAM window displays.

5. Click **Browse** to select the NVSRAM file that you previously downloaded to a directory on the host.

6. Click **OK**. The Confirm Download window displays.

7. Click **Yes** to start the download.

8. Verify that the NVSRAM was successfully installed.

9. Start I/O.

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 24.

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"

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 Topology section of the Mappings window, highlight the name of the storage subsystem, and click **Mappings —> Define —> Host Group**.

   Note: Make sure that the storage subsystem is highlighted in the left panel of the Subsystem Management window. Do not highlight Undefined Mappings.
3. Type the name of the new host group (for example, AIX). Click **Add**, and then click **Close**.
4. In the Topology section of the Mappings window, 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

This section contains the following procedures:

- "Mapping LUNs to a new partition"
- "Adding LUNs to an existing partition"

Mapping LUNs to a new partition

When mapping LUNs for the first time to a newly created partition, complete the following steps to map LUNs to the partition.

Before you begin: Read "DS4000 Storage Manager performance planning" on page 21 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 section, right-click the host on which you want to map LUNs, and select **Define Storage Partitioning**. The Define Storage Partitioning window opens.
3. In the Define Storage Partitioning window, select **Host**, then click **Next**.
4. Select the logical drive by name, on the right side of the window.
5. Accept the default LUN ID, or change it, then click **Add**.
6. Repeat step 5 for each LUN that you want to map to the partition.

Adding LUNs to an existing partition

When adding LUNs to an existing partition, complete the following steps to map the new LUNs to the partition.

Before you begin: Read "DS4000 Storage Manager performance planning" on page 21 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 section, 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 that you want to add to the partition.

Creating direct and SAN-attached configurations

DS4000 Storage Manager supports IBM DS4000 Fibre Channel Storage Subsystems 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 subsystems.
- 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 subsystems 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 subsystem.
2. Configure and verify the configuration, using the procedures that are described in "Identifying devices on AIX hosts" on page 44 and "Verifying the installation and configuration" on page 46.

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 F, "Connecting HBAs in a FC switch environment," on page 223.

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 F, "Connecting HBAs in a FC switch environment," on page 223.
4. Configure and verify the configuration.
5. Use the **cfgmgr** command to enable the HBAs 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 44 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 43, Table 12 on page 43, and Table 13 on page 43 must be the specified versions or later. Be sure to check the readme file for more recently-updated information.

A Storage Manager readme file for AIX, which 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 subsystem. When the storage subsystem page opens, click the **Download** tab.

You can also find controller firmware versions listed in Table 3 on page 1.

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 42.

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	5.1.0.65
devices.common.IBM.fc.rte	5.1.0.51
devices.pci.df1000f7.com	5.1.0.64
devices.pci.df1000f7.rte	5.1.0.37
devices.pci.df1000f9.rte	5.1.0.37

Table 12. File sets required for AIX 5.2 RDAC

PTF file sets	Version
devices.fcp.disk.array.diag	5.2.0.30
devices.fcp.disk.array.rte	5.2.0.60
devices.common.IBM.fc.rte	5.2.0.50
devices.pci.df1000f7.com	5.2.0.60
devices.pci.df1000f7.rte	5.2.0.30
devices.pci.df1000f9.rte	5.2.0.30
devices.pci.df1000fa.rte	5.2.0.50

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	5.3.0.20
devices.common.IBM.fc.rte	5.3.0.0
devices.pci.df1000f7.com	5.3.0.21
devices.pci.df1000f7.rte	5.3.0.0
devices.pci.df1000f9.rte	5.3.0.0
devices.pci.df1000fa.rte	5.3.0.10

Important: With AIX 5.3, download the complete maintenance package and update all PTFs together. Do not install each PTF separately.

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 29 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:

```
# lslpp -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.

3. 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 on page 45:

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.
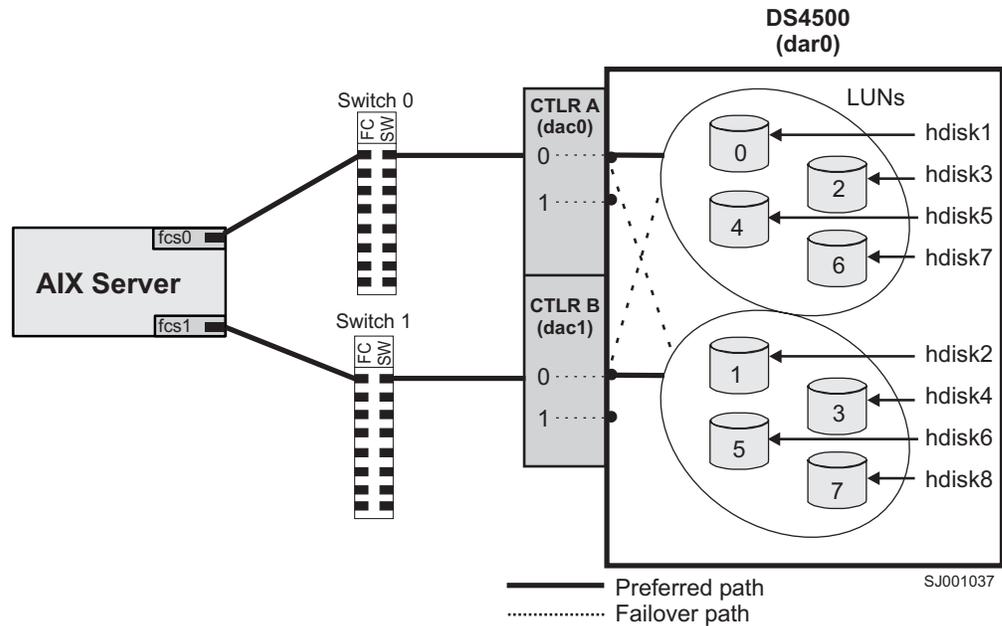
Figure 8. DS4000/AIX system configuration

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".
   - Each DS4800 logical drive is recognized as a "1815 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 55.

See "Using the lsattr command to view ODM attributes" on page 61 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 48
- "Identifying logical drives by operating system device names" on page 53
- "Identifying FlashCopy logical drive types" on page 54

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" on page 47

- "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 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 on page 48, 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

As mentioned previously in this chapter, the RDAC software uses *dars* to represent the disk array routers, *dacs* to represent the disk array controllers, and *hdisks* to represent the logical drives.

After these operating system device names are found, you can use one or more of the following methods in AIX to get information about device names, bus numbers, and other attributes. You can then use this information to correlate the device names with their logical drives and preferred and alternate paths:

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

"Using the fget_config command" on page 50
> The **fget_config** command displays controllers and hdisks that are associated with a specified DS4000 (dar). **fget_config** also shows the current state and logical drive (hdisk) ownership.

"Using the lsattr command" on page 52
> The **lsattr** command displays device attributes and possible values. Attributes are only updated at startup time, or the last time that the **cfgmgr -v** command was run.

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-adapter 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 -vl dar3

User array name = 'D59 DS4500 #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---                              ---dar2---

dac3 ACTIVE dac0 PASSIVE                dac2 ACTIVE dac5 ACTIVE
hdisk4    dac3                          hdisk25   dac5
hdisk5    dac3
hdisk6    dac3                          ---dar3---
hdisk7    dac3
hdisk8    dac3                          dac6 ACTIVE dac9 ACTIVE
hdisk10   dac3                          hdisk26   dac9
hdisk11   dac3                          hdisk27   dac6
hdisk12   dac3                          hdisk28   dac9
hdisk13   dac3                          hdisk29   dac6
hdisk14   dac3                          hdisk30   dac9
                                        hdisk31   dac6
---dar1---                              hdisk32   dac9
                                        hdisk33   dac9
dac1 ACTIVE dac4 ACTIVE                 hdisk83   dac6
hdisk15   dac1
hdisk16   dac1
hdisk17   dac4
hdisk18   dac1
hdisk19   dac4
hdisk20   dac1
hdisk21   dac4
hdisk22   dac1
hdisk23   dac4
hdisk24   dac1
```

_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---                                      ---dar2---

User array name = 'D59S Flute-2              User array name = 'D59S DS4500 #2'
(DS4500) #1'                                 dac2 ACTIVE dac5 ACTIVE
dac3 ACTIVE dac0 PASSIVE
                                             Disk      DAC    LUN Logical Drive
Disk      DAC    LUN Logical Drive           utm               31
hdisk4    dac3    0 dualpart0                hdisk25   dac5    0 1
hdisk5    dac3    1 dualpart1
hdisk6    dac3    2 dualpart2                ---dar3---
hdisk7    dac3    4 dualpart4
hdisk8    dac3    5 dualpart5                User array name = 'D59 DS4500 #1'
hdisk10   dac3    8 dualpart8                dac6 ACTIVE dac9 ACTIVE
hdisk11   dac3    9 dualpart9
hdisk12   dac3   10 dualpart10               Disk      DAC    LUN Logical Drive
hdisk13   dac3   11 dualpart11               utm               10
hdisk14   dac3   12 dualpart12               hdisk26   dac9    1 multipart11
                                             hdisk27   dac6    2 multipart12
---dar1---                                   hdisk28   dac9    3 multipart13
                                             hdisk29   dac6    4 multipart14
User array name = 'D59 DS4500 #1'            hdisk30   dac9    5 multipart15
dac1 ACTIVE dac4 ACTIVE                      hdisk31   dac6    6 multipart16
                                             hdisk32   dac9    7 multipart17
Disk      DAC    LUN Logical Drive           hdisk33   dac9    9 multipart19
utm               10                         hdisk83   dac6    8 multipart18
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
```

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 600A0B80000F14AF000003D140C046A2.

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**.
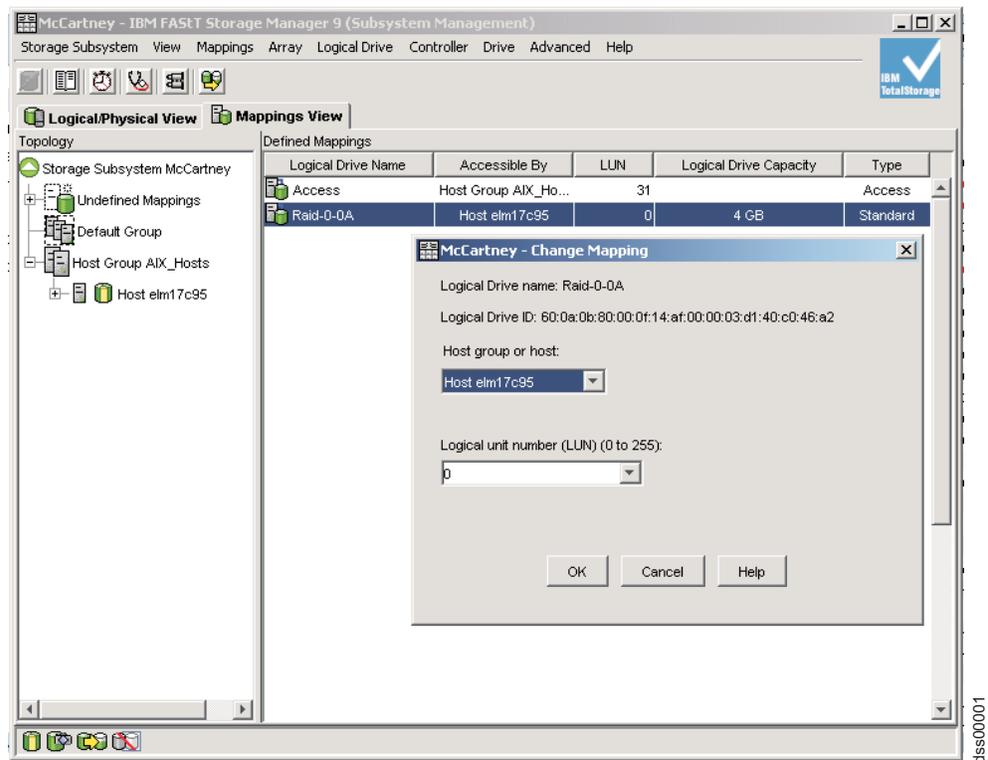
Figure 10. Change Mappings window

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 48.

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 subsystem 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 56: *Attributes for dar devices*
- Table 15 on page 57: *Attributes for dac devices*
- Table 16 on page 58: *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 61 for more information.

Table 14. Attributes for dar devices

Attribute	Definition	Changeable (T/F)	Possible value
act_controller	List of controllers in the active state at the time of configuration.	False	Set at configuration time by the RDAC software.
all_controller	List of controllers that comprise this array; usually there are two dac devices.	False	Set at configuration time by the RDAC software.
held_in_reset	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.
load_balancing	Indicator that shows whether load balancing is enabled (**yes**) or disabled (**no**); see the definition of the *balance_freq* attribute for more information.	True	Yes or No. **Attention:** You should only set the *load_balancing* attribute to **yes** in single-host configurations.
autorecovery	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.
hlthchk_freq	Number that specifies how often health checks are performed, in seconds.	True	1 - 9999. Should not be changed
aen_freq	Number that specifies how often polled AEN checks are performed, in seconds.	True	1 - 9999. Should not be changed
balance_freq	If *load_balancing* 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
fast_write_ok	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.
cache_size	Cache size for both controllers, in megabytes; **0** if the sizes do not match.	False	512 or 1024. Set by DS4000.
switch_retries	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
passive_control	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.
alt_held_reset	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.
controller_SN	Serial number of this controller.	False	Set by DS4000.
ctrl_type	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
cache_size	Cache size of this controller, in megabytes.	False	512, 1024. Set by DS4000.
scsi_id	SCSI identifier of this controller.	False	Set by SAN, reported by AIX.
lun_id	Logical unit number of this controller.	False	Set by DS4000.
utm_lun_id	Logical unit number of this controller, or **none** if UTM (access logical drives) is not enabled.	False	0 - 31. Set by DS4000 Storage Manager.
node_name	Name of the Fibre Channel node.	False	Set by DS4000
location	User-defined location label for this controller; the system does not use this value.	True	Set by DS4000 Storage Manager.
ww_name	Fibre Channel worldwide name of this controller.	False	Set by DS4000.
GLM_type	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
pvid	AIX physical volume identifier, or **none** if not set.	False	Set by AIX.
q_type	Queueing type for this device; must be set to **simple**.	False	Set by AIX. Must be "simple".
queue_depth	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 60 for important information about setting this attribute.
PR_key_value	Required only if the device supports any of the persistent reserve policies. This attribute is used to distinguish between different hosts.	True	1-64, or None. **Note:** You must set this attribute to non-zero before the reserve_policy attribute is set.

Table 16. Attributes for hdisk devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
reserve_policy	Persistent reserve policy, which defines whether a reservation methodology is employed when the device is opened.	True	no_reserve PR_shared, PR_exclusive, or single_path_reserve.
max_transfer	Maximum transfer size is the largest transfer size that can be used in sending I/O.	True	Numeric value; Default = 1 MB **Note:** Usually unnecessary to change default, unless very large I/Os require increasing the value.
write_cache	Indicator that shows whether write-caching is enabled on this device (**yes**) or not (**no**); see the definition of the *cache_method* attribute for more information.	False	Yes or No.
size	Size of this logical drive.	False	Set by DS4000.
raid_level	Number that specifies the RAID level of this device.	False	0, 1, 3, 5. Set by DS4000 Storage Manager.
rw_timeout	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.
reassign_to	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.
scsi_id	SCSI identifier at the time of configuration.	False	Set by SAN, reported by AIX.
lun_id	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
cache_method	If *write_cache* is enabled, the write-caching method of this array; set to one of the following: • **default**. Default mode; the word ″default″ is not seen if *write_cache* 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.
prefetch_mult	Number of blocks to be prefetched into read cache for each block read.	False	0 - 100.
ieee_volname	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 performance.

This section provides methods for calculating your system's maximum queue depth, which you can use as a guideline to help you determine the best queue depth setting for your configuration.

Calculating maximum queue depth: The formula for calculating the maximum queue depth for your system depends on which firmware version is installed on the controller. Use one of the following formulas to calculate the maximum queue depth for your system.

Important:

1. The maximum queue depth might not be an optimal setting in all cases. Use the maximum queue depth as a guideline, and adjust the setting as necessary for your specific configuration.
2. In systems with one or more SATA devices attached, you might need to set the queue depth attribute to a lower value than the maximum queue depth.

- **Formula for controller firmware versions 06.1x.xx.xx or 05.4x.xx.xx**

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

  ```
  2048 / (number-of-hosts * 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
  ```

- **Formula for controller firmware version 05.30.xx.xx**

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

  ```
  512 / (number-of-hosts * 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
  ```

Setting queue depth attribute: You can set the *queue_depth* attribute using the **chdev -l** command, as shown in the following example:

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

where **X** is the name of the hdisk and **y** is the queue depth setting.

Note: Use the **-P** flag to make the changes permanent in the Customized Devices object class.

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 -Dl**.
- To view the attributes that are currently set on the system, type **lsattr -El**.

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

Example 1: Attribute settings for dar0:

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

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 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" on page 63
- "Using fast I/O failure for Fibre Channel devices" on page 63
- "Using dynamic tracking of Fibre Channel devices" on page 64
- "Using dynamic capacity expansion and dynamic volume expansion" on page 65

Disabling DS4000 cache mirroring

Cache mirroring is enabled on the DS4000 storage subsystem 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 Logical/Physical view of 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 72.

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 subsystem. 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 subsystem. 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 subsystem.
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 subsystem.
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. In the Logical/Physical view of the Subsystem Management window, right-click the logical drive whose capacity you want to increase, and select **Increase Capacity**. The Increase Logical Drive Capacity — Additional Instructions window opens.
2. Read the additional instructions and click **OK**. The Increase Logical Drive Capacity window opens.
3. 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.

Replacing hot swap HBAs

This section describes the procedure for hot-swapping Fibre Channel host bus adapters (HBAs) on an AIX host.

The hot swap procedure requires knowledge in the following areas:

- AIX administration
- AIX hardware procedures for PCI card replacement.
- IBM TotalStorage DS4000 Storage Manager
- Fibre Channel zoning procedures (required if a switch is used to interconnect DS4000 storage subsystem and the AIX host, and zoning is based on WWPN)

> **Caution:** If the procedures are not followed as they are documented in this section, loss of data availability can occur. IBM recommends that you read and understand all of the requirements and steps in this section before you begin the HBA hot swap procedure.

This section contains the following information:

- "Requirements"
- "Known issues and restrictions" on page 67
- "Preparing for the HBA hot swap" on page 67
- "Replacing the hot swap HBA" on page 70
- "Mapping the new WWPN to the DS4000 storage subsystem" on page 71
- "Completing the HBA hot swap procedure" on page 71

Requirements

The following are the minimum AIX operating system and file set versions that are required to support hot swap:

Table 17. Required AIX operating systems and file set versions for supporting HBA hot swap

Operating system	Package name	File set version
AIX 5.1	devices.fcp.disk.array.rte	5.1.0.65
AIX 5.2	devices.fcp.disk.array.rte	5.2.0.60
AIX 5.3	devices.fcp.disk.array.rte	5.3.0.20

Known issues and restrictions

Please note the following known issues and restrictions when you perform a hot swap operation:

Caution: Any deviations from these notes and procedures might cause a loss of data availability.

- The autorecovery attribute of the dar must be set to `no`. Autorecovery is a dynamically set feature that can be turned back on after the hot swap procedure is complete. Failure to disable autorecovery mode during a hot swap procedure can cause loss of access to data.
- Do not redistribute logical drives to the preferred path until you verify that the HBA replacement succeeded and that the subsequent configuration was performed correctly. Redistributing the logical drives before verifying successful hot swap and configuration can cause a loss of access to data.
- The only supported hot swap scenario is the following operation:
  - Replacing a defective HBA with the *same* model HBA, and in the *same* PCI slot.

  Do not insert the defective HBA into any other system, even if the HBA is found not to actually be defective. Always return the HBA to IBM.

  Important: No other variations of replacement scenarios are currently supported.
- Hot swap is not supported in single-HBA configurations.

Preparing for the HBA hot swap

Complete the procedures in this section to prepare for the hot swap.

Collecting system data

In preparation for the hot swap procedure, complete the following steps to collect data from the system:

1. Type the following command:

```
# lsdev -C |grep fcs
```

   The output is similar to the following example:

```
fcs0        Available 17-08       FC Adapter
fcs1        Available 1A-08       FC Adapter
```

2. Type the following command:

```
# lsdev -C |grep dac
```

   The output is similar to the following example:

```
dac0        Available 17-08-02    1815    DS4800 Disk Array Controller
dac1        Available 1A-08-02    1815    DS4800 Disk Array Controller
```

3. Type the following command for each of the fcs devices:

```
# lscfg -vpl fcsX
```

   where *X* is the number of the fcs device. The output looks similar to the following example:

```
lscfg -vpl fcs0
 fcs0            U0.1-P1-I1/Q1  FC Adapter

        Part Number.................09P5079
        EC Level....................A
        Serial Number...............1C21908D10
        Manufacturer................001C
        Feature Code/Marketing ID...2765
        FRU Number..................09P5080
        Network Address.............10000000C92D2981
        ROS Level and ID............02C03951
        Device Specific.(Z0)........2002606D
        Device Specific.(Z1)........00000000
        Device Specific.(Z2)........00000000
        Device Specific.(Z3)........03000909
        Device Specific.(Z4)........FF401210
        Device Specific.(Z5)........02C03951
        Device Specific.(Z6)........06433951
        Device Specific.(Z7)........07433951
        Device Specific.(Z8)........20000000C92D2981
        Device Specific.(Z9)........CS3.91A1
        Device Specific.(ZA)........C1D3.91A1
        Device Specific.(ZB)........C2D3.91A1
        Device Specific.(YL)........U0.1-P1-I1/Q1


   PLATFORM SPECIFIC

     Name:  fibre-channel
     Model:  LP9002
     Node:  fibre-channel@1
     Device Type:  fcp
     Physical Location: U0.1-P1-I1/Q1
```

4. Type the following command:

```
# lsdev -C |grep ^dar
```

The output looks similar to the following example:

```
# dar0        Available              1815     DS4800 Disk Array Router
  dar1        Available              1815     DS4800 Disk Array Router
```

5. Type the following command to list the attributes of each dar found on the system:

```
# lsattr -El darX
```

where *X* is the number of the dar. The output looks similar to the following example:

```
lsattr -El dar0
 act_controller  dac0,dac2  Active Controllers                            False
 all_controller  dac0,dac2  Available Controllers                         False
 held_in_reset   none       Held-in-reset controller                      True
 load_balancing  no         Dynamic Load Balancing                        True
 autorecovery    no         Autorecover after failure is corrected        True
 hlthchk_freq    600        Health check frequency in seconds             True
 aen_freq        600        Polled AEN frequency in seconds               True
 balance_freq    600        Dynamic Load Balancing frequency in seconds   True
 fast_write_ok   yes        Fast Write available                          False
 cache_size      1024       Cache size for both controllers               False
 switch_retries  5          Number of times to retry failed switches      True
```

Verifying that autorecovery is disabled

Before you perform the hot swap, you must complete the following steps to ensure
that autorecovery is disabled on every dar that is involved with the HBA you want to
hot swap:

1. Identify all the dac(s) that are involved with the HBA by typing the following
    command:

```
# lsdev -C|grep 11-08
```

The output looks similar to the following example:

```
# lsdev -C|grep 11-08
fcs0      Available 11-08          FC Adapter
fscsi0    Available 11-08-01       FC SCSI I/O Controller Protocol Device
dac0      Available 11-08-01       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
```

2. Consult the **lsattr** command output that you collected in step 5 on page 68 of
    the procedure "Collecting system data" on page 67. In the **lsattr** output, identify
    the dar(s) that list the dacs you identified in step 1 of this procedure.

3. For each dar that you identified in step 2, type the following command:

```
#  lsattr -El darX |grep autorecovery
```

where *X* is the number of the dar. The output looks similar to the following
example:

```
#  lsattr -El dar0 |grep autorecovery
autorecovery   no        Autorecover after failure is corrected       True
```

4. In the lsattr command output, verify that the second word is **no**. If the second
    word is set to **yes**, then autorecovery is currently enabled.

    Important: For each dar on which autorecovery is enabled, you must disable it
    by setting the autorecovery ODM attribute to **no**. See "Viewing and
    setting Object Data Manager (ODM) attributes" on page 55 to learn
    how to change attribute settings. Do not proceed with the hot swap
    procedure until you complete this step and verify that autorecovery
    is disabled.

Replacing the hot swap HBA

After you complete the procedures in "Preparing for the HBA hot swap" on page 67, complete the following steps to replace the hot swap HBA:

1. Place the HBA that you want to replace into the `Defined` state by typing the following command:

```
# rmdev -Rl fcsX
```

   where *X* is the number of the HBA. The output is similar to the following example:

```
rmdev -Rl fcs0
   fcnet0 Defined
   dac0 Defined
   fscsi0 Defined
   fcs0 Defined
```

2. In the AIX smit menu, initiate the process that is required for the HBA hot swap by selecting **smit --> Devices --> PCI Hot Plug Manager --> Replace/Remove a PCI Hot Plug Adapter**.

3. In the Replace/Remove a PCI Hot Plug Adapter window, select targeted HBA. A window displays that contains instructions for replacing the HBA.

4. Replace the HBA by following the smit instructions.

   Note: Do not reinstall the Fibre Channel Cable at this time.

5. If the steps in this procedure are completed successfully up to this point, you obtain the following results:

   - The "defective" HBA is removed from the system.
   - The replacement FC HBA is powered on.
   - The associated fcs*X* device is in the `Defined` state.

   Before continuing, verify that these results have been obtained.

6. Install the Fibre Channel loop back on the replacement HBA.

7. Place the HBA into the `Active` state by typing the following command:

```
# cfgmgr
```

8. Verify that the fcs device is now available by typing the following command:

```
# lsdev -C |grep fcs
```

9. Verify or upgrade the firmware on the replacement HBA to the appropriate level by typing the following command:

```
# lscfg -vpl fcsX
```

   where *X* is the number of the fcs.

   Note: You can determine the HBA firmware level by referring to the fcs*X* device data that you gathered during the procedure at the start of this section, "Collecting system data" on page 67.

10. Record the 16-digit number that is associated with `Network Address`, as it was displayed in the output of the command you used in step 9. This Network

Address number will be used in the next procedure, manually map the
replacement HBA's WWPN to the Storage subsystem(s).

11. Place the HBA back into the `Defined` state by typing the following command:

```
# rmdev -Rl fcsX
```

When you have completed this procedure, continue to the next procedure,
"Mapping the new WWPN to the DS4000 storage subsystem."

Mapping the new WWPN to the DS4000 storage subsystem

For each DS4000 storage subsystem that is affected by the hot swap, complete the
following steps to map the worldwide port name (WWPN) of the HBA to the storage
subsystem:

1. Start DS4000 Storage Manager and open the Subsystem Management window.

2. In the Mapping View of the Subsystem Management window, select **Mappings
   --> Show All Host Port Information**. The Host Port Information window
   displays.

3. Using the data that you collected during the procedure "Collecting system data"
   on page 67, find the entry in the Host Port Information window that matches the
   WWPN of the "defective" HBA (the HBA that you removed), and record the alias
   name. Then, close the Host Port Information window.

4. In the Mapping View, select the alias name of the HBA host port that you just
   recorded.

5. Select **Mappings --> Replace Host Port**. The Replace Host Port window
   opens.

6. In the Replace Host Port window, verify that the current HBA Host Port
   Identifier, which is listed at the top of the window, exactly matches the WWPN of
   the HBA that you removed.

7. Type the 16-digit WWPN of the replacement HBA in the New Identifier field, and
   click **OK**.

When you have completed these steps continue to the next procedure, "Completing
the HBA hot swap procedure."

Completing the HBA hot swap procedure

Complete the following steps to finish replacing the hot swap HBA:

1. Remove the Fibre Channel loop back plug, and insert the Fibre Channel cable
   that was previously attached to the HBA that you removed.

2. If HBA is attached to a Fibre Channel switch, *and* the zoning is based on
   WWPN, modify the zoning information to replace the WWPN of the former HBA
   with the WWPN of the replacement HBA. (Run **cfgmgr** at this time to allow the
   HBA to register its WWPN in the Fibre Channel switch.)

   Important: Skip this step if the HBA is directly attached to the DS4000
   subsystem, or if the Fibre Channel switch zoning is based on port
   numbers instead of WWPNs. If you do need to modify the zoning,
   failure to correctly do so will prevent the HBA from accessing the
   storage subsystem.

3. Run the **cfgmgr** command.

4. Type the following commands to verify that the replaced fcsX device and its
   associated dac(s) are placed in the `Available` state:

```
# lsdev -C |grep fcs

lsdev -C |grep dac
```

5. Type the following step to verify that no additional dar(s) have been created and that the expected dar(s) are in the `Available` state. (Refer to the data that you collected during the procedure "Collecting system data" on page 67 to compare the original number of dar(s) to the number that is now reported by the system.)

```
# lsdev -C |grep dar
```

> **Caution:** The presence of additional dar(s) in the **lsdev** output indicates a configuration problem. If this occurs, do *not* continue this procedure until you correct the problem, Loss of data availability can occur.

6. For each dar, type the following command to verify that affected dar attributes indicate the presence of two active dac(s):

```
# lsattr -El darX|grep act_controller
```

where *X* is the number of the dar. The output looks similar to the following:

```
lsattr -El dar0|grep act_controller
act_controller dac0,dac2 Active Controllers                         False
```

> **Caution:** If two dacs are not reported for each affected dar, loss of data availability can occur. Do *not* continue this procedure if two dac(s) are not reported for each dar. Correct the problem before continuing.

7. Using the Storage Manager procedure that is described in "Redistributing logical drives in case of failure" on page 75, manually redistribute volumes to preferred paths.

8. Verify that disks stay on preferred path by using one or both of the following methods:

   Using AIX system
   > Run the **fget_config -Av** command, and verify that drives are on expected path

   Using Storage Manager
   > In the Enterprise Management window, verify that the storage subsystem(s) are `Optimal`. If they are not `Optimal`, verify that any drives that are part of the subsystems involved with hot swap process are *not* listed in the Recovery GURU.

9. Enable autorecovery of the affected dar(s) at this time, if needed. (See "Viewing and setting Object Data Manager (ODM) attributes" on page 55 to learn how to change attribute settings.)

Result: The Fibre Channel HBA hot swap is now complete.

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.
- **FCP_ARRAY_ERR27** *SINGLE CONTROLLER RESTARTED*

  The subsystem is operating as a single controller, and an error has been repaired. The error might have been a communication or hardware problem, or it might have occurred because a LUN was moved to a controller that does not have a path to the current host

  If this is a dual-controller subsystem, find the reason that the subsystem is operating in single-controller mode, and resolve the problem. Possible reasons include:
  - – An HBA, switch port, switch, DS4000 port or DS4000 controller was unavailable during the last system reboot or the last time the **cfgmgr** command was run.
  - – A user removed a path (dac) as part of a Fibre Channel adapter hot swap operation.
- **FCP_ARRAY_ERR28** *SINGLE CONTROLLER RESTART FAILURE*

  The subsystem is operating as a single controller, and the error has not been repaired. There is a problem with the path between this host and the subsystem or with the subsystem itself. The host has attempted to communicate with the subsystem and that communication has failed.

  If the number of retries that is specified in the ODM attribute *switch_retries* is reached, the I/O is failed back to the user.

  Repair the error. Then, if this is a dual-controller subsystem, find the reason that the subsystem is operating in single-controller mode, and resolve that problem. Possible reasons include:
  - – An HBA, switch port, switch, DS4000 port or DS4000 controller was unavailable during the last system reboot or the last time the **cfgmgr** command was run.
  - – A user removed a path (dac) as part of a Fibre Channel adapter hot swap operation.

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 host 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 subsystem.
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 80
- "Configuring storage subsystems" on page 85
- "Creating direct-attached and SAN-attached configurations" on page 91
- "Enabling multipath I/O with PV-links" on page 92
- "Redistributing logical drives in case of failure" on page 97

Hardware, firmware, and host system requirements

This section lists 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 Storage Manager readme file for HP-UX, 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 subsystem. When the storage subsystem page opens, click the **Download** tab, and follow the links to the readmes.

DS4000 hardware and firmware requirements

Table 18 lists the DS4000 Fibre Channel storage subsystems that you can use with DS4000 Storage Manager.

Table 18. DS4000 storage subsystems that you can use with HP-UX systems

Product name	Model	Controller firmware versions
IBM TotalStorage DS4800 Storage Subsystem	1815-82A/H, 1815-84A/H	06.14.xx.xx
IBM TotalStorage DS4100 SATA Base Storage Subsystem	1724-100	06.12.xx.xx
IBM TotalStorage DS4100 SATA Single-controller (SCU) Storage Subsystem	1724-1SX	05.42.xx.xx
IBM TotalStorage DS4500 Storage Subsystem	1742-90U	06.12.xx.xx
IBM TotalStorage DS4400 Storage Subsystem	1742-1RU	06.12.xx.xx
IBM TotalStorage DS4300 Standard and Turbo Storage Subsystem	1722-60U, 1722-60X	06.12.xx.xx
IBM TotalStorage DS4300 single-controller (SCU) Storage Subsystem	1722-6LU, 1722-6LX	5.34.xx.xx

Table 18. DS4000 storage subsystems that you can use with HP-UX systems (continued)

Product name	Model	Controller firmware versions
NVSRAM: For the latest NVSRAM versions, see the most recent Storage Manager readme files for HP-UX at the following Web site: www-1.ibm.com/servers/storage/support/disk/ Click the link for your storage subsystem. When the storage subsystem page opens, click the **Download** tab, then click the Storage Manager link, then click the **Storage Mgr** tab.		

Table 19 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 19. 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 **Note:** When upgrading the EXP100 to ESM firmware version 9554, you must use a storage subsystem that has DS4000 controller firmware **05.41.1x.xx** installed . You cannot upgrade to ESM firmware 9554 using a storage subsystem with different firmware installed.
	IBM TotalStorage DS4000 EXP710 Fibre Channel Storage Expansion Enclosure	1740-710	9627
	IBM TotalStorage DS4000 EXP700 Fibre Channel Storage Expansion Enclosure	1740-1RU or 1RX	9326 **Note:** You must upgrade the EXP700 ESM firmware to version 9326 before you upgrade the DS4000 controller firmware to version **06.1x.xx.xx**.
	IBM TotalStorage DS4000 EXP500 Fibre Channel Storage Expansion Enclosure	3560-1RU or 1RX	9166
Host Bus Adapters	HP FC Adapter for HP-UX 11.0	A5158A, A6795A	B.11.00.10,
	HP FC Adapter for HP-UX 11i	A5158A, A6795A	B.11.11.09, B.11.11.03

Table 19. Supported versions of hardware for HP-UX systems (continued)

Hardware	Product Name	Model	Minimum firmware version
Switches	IBM TotalStorage SAN Switch	2109-S08 2109-S16	2.6.2A
	IBM TotalStorage SAN Switch	2109-F16	3.1.2A
	IBM TotalStorage SAN switch	2109-M12	4.0.2C
	IBM TotalStorage SAN switch	2109-F32	4.2.0C
	IBM TotalStorage SAN 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
	CNT switch	2042-001	4.3.0
	Cisco switch	2062-D01, 2062-D07, 2062-T07	1.3.4A

Note: If you have any questions about the latest DS4000 interoperability information, see the DS4000 interoperability matrix at the following Web site:

www-1.ibm.com/servers/storage/disk/ds4000/interop-matrix.html

HP-UX system requirements

This section lists the minimum hardware and software requirements for an HP-UX host in a DS4300, DS4400, DS4500, or DS4800 storage subsystem that is managed by DS4000 Storage Manager 9.1x.

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 SMutil and SMagent
- 47.5 MB available disk space on /opt and at least 95 MB available on /tmp for installation in an SMruntime environment

Software requirements

Ensure that the HP-UX host is running one of the following operating systems with the appropriate patches. Because 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, at minimum:
  - PHKL_23939
  - PHKL-27003

                – PHKL_24027

                – PHKL_27364

                – 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, at minimum:

                – PHKL_26519

                – PHKL_26743

                – 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 20.

Table 20. 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

The DS4000 Storage Manager software includes the following software packages:

SMruntime (required)
> SMruntime is a Java compiler for the DS4000 Storage Manager software packages. SMruntime must be installed before the other packages are installed.

SMclient (optional)
> SMclient is Storage Manager's graphical user interface (GUI), which you use to manage your storage subsystems. If you are not using your HP-UX host as a storage management station (for example, if you are using a Windows management station with SMclient installed), then you do not need to install SMclient on your HP-UX host.

SMagent (optional)
> SMagent contains host-agent software, which you can use to manage your subsystems if you are using an in-band management configuration (managing the storage subsystems through the Fibre Channel I/O path to the host). If you are using an out-of-band management configuration, you do not need to install SMagent.

SMutil (required)
> SMutil is a utility that you can use to register and map new logical drives to the operating system. SMutil is required on HP-UX host systems.

You can install all of the DS4000 Storage Manager software packages automatically by using the Storage Manager installation wizard, or you can install each package manually. This section contains the following procedures:

      •

Installing Storage Manager using the installation wizard

The DS4000 Storage Manager installation wizard can install all of the Storage Manager software packages automatically. If you prefer, you can install just the following required packages onto your HP-UX host:

• SMruntime

• SMutil

To choose the packages that you want to install, select Custom Installation during the installation procedure, as described in "Steps for installing Storage Manager using the installation wizard."

Installation wizard requirements

To use the installation wizard, you must meet the following requirements:

• You are installing DS4000 Storage Manager 9.12 or 9.14

• Your HP-UX host has a graphics adapter card installed.

If you are installing Storage Manager 9.10 or an earlier version, or if your system does not have a graphics card installed, skip this section and install the stand-alone host software packages using the procedures described in "Installing Storage Manager packages manually." All the packages are included with the installation CD.

Note: Do not install the software packages in the same server by using both the installation wizard and the manual installation method.

Steps for installing Storage Manager using the installation wizard

Complete the following steps to install the DS4000 Storage Manager software using the installation wizard. Adjust the steps as necessary for your specific installation.

1. Download the following file from the DS4000 Storage Manager CD, or from the DS4000 support Web site, to a directory on your system:

   SMIA-HPUX-09.1x.xx.xx.bin

2. Open the file on your system. The Storage Manager Installation wizard's Introduction window opens.

3. Follow the instructions in each window of the wizard. When you select the installation type, you can choose one of the following options:

   • Typical (Full) Installation — Installs all Storage Manager software packages

   • Management Station — Installs SMruntime and SMclient

   • Host — Installs SMruntime, SMagent and SMutil

   • Custom — Allows you to select which packages you want to install

4. Restart the system.

Result: The DS4000 Storage Manager software is installed on your system.

Installing Storage Manager packages manually

Use the following procedures, in the order listed, to install DS4000 Storage Manager 9.10 (or earlier) 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" on page 82

2. "Installing the host software" on page 83

Note: If you are installing Storage Manager 9.12 or 9.14, do not install the software manually. Instead, use the installation procedure that is described in "Installing Storage Manager using the installation wizard" on page 81.

Installing the client software

Install the client software in the following order:

1. DS4000 SMruntime software (required)
2. DS4000 SMclient software (optional)

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 79.
- 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 (required):   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_09xxxxx
```

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 (optional): If you are using your HP-UX host as a storage management station, then you must install SMclient on the host. If you are using a different computer as your storage management station, install SMclient on that computer.

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-09xxxxx
```

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 79.
- SMruntime has been installed.

Note: SMagent is not supported with DS4000 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-09xxxxxx
```

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_09xxxxxx
```

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 179 for additional information about storage configuration.

Before you begin: Note the following information:
- All DS4000 storage subsystems 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 controller 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 Storage Manager readme file for HP-UX 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 subsystem. When the storage subsystem page opens, click the **Download** tab. Follow the links to the readmes.

   Note: If you do not have the latest versions of controller firmware and NVSRAM, download them and complete the procedure described in "Upgrading controller firmware and NVSRAM" on page 86. 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 26 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 Storage Subsystem**.

   b. In the Add Storage Subsystem window, type the IP addresses of each controller in the storage subsystem, and click **Add**. The name of the storage subsystem displays in the Enterprise Management window.

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

4. In the Enterprise Management window, double-click the name of the storage subsystem to open its Subsystem Management window.

5. 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.

Overview

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

- "Using concurrent download"
- "Using traditional or staged download"

Using concurrent download: Depending on your system's current firmware, 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.

Using traditional or staged download: There are two methods for downloading the firmware to the controllers:

- "Traditional download"
- "Staged download" on page 87

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 DS4000 controller firmware version 5.4x.xx.xx to version 6.1x.xx.xx. You can only use staged download if the DS4000 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 subsystems.

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 subsystem 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 subsystem that you plan to update.

Preparing for firmware or NVSRAM download

To prepare for download, review and complete all of the steps in this section before you start the download procedures.

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

To check the controller 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 view the current versions of controller firmware and NVSRAM.
3. Click **Close** to close the Storage Subsystem Profile window.
4. To find whether the versions of controller firmware and NVSRAM that are currently on your storage subsystem are the latest versions, see the Storage Manager readme file for HP-UX at the following Web site:

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

   Click the link for your storage subsystem.

   When the storage subsystem page opens, click the **Download** tab, then click the Storage Manager tab.

   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, download the correct versions from this Web site to a directory on your system, then follow the preparation and upgrade procedures that are described in this section.

Save the storage subsystem configuration: Saving the storage subsystem configuration to a script file is necessary in order to restore the configuration in the event of a catastrophic failure.

Attention: Do not use this option if the storage subsystem is undergoing an operation that changes any configuration settings. (If any logical drive listed in the Logical View of the Subsystem Management window has a clock icon, then a configuration change is occurring, and you must wait to save the configuration.)

You can also save the storage subsystem profile, by selecting **Storage Subsystem —> View Profile** in the Subsystem Management window, and clicking **Save As**. Then, in the Save Profile window, select the sections of the profile that you want to save, specify a name and location for the text file, and click **Save**.

Save the storage subsystem configuration by completing the following steps:

1. In the Subsystem Management window, click **Storage Subsystem —> Configuration —> Save**. The Save Configuration window displays.
2. In the Save Configuration window, select the elements of the configuration that you want to save, and click **Yes**. A save dialog displays.
3. In the save dialog, specify a name for the script file and select its location.

   Notes:

   - The .cfg extension is used by default for the script file if you do not specify a file extension.
   - Do not store your DS4000 Storage Subsystem configuration script file in the same location as the logical drives that are defined on your DS4000 Storage Subsystem. If you do, and the logical drives fail during the subsequent migration, you will lose data.

4. Click **Save**. The Configuration Saved window displays.
5. Review the information in the Configuration Saved window, then click **OK** to return to the Subsystem Management window.

Determine whether you need to stop I/O: You can use concurrent download only if you are upgrading from one version of controller firmware to a later release of the same version. For example, if you are upgrading from 06.10.xx.xx to 6.12.xx.xx, you can use concurrent download. However, if you are upgrading from 5.4x.xx.xx to 06.1x.xx.xx, you *cannot* use concurrent download.

In addition, before you upgrade controller firmware, ensure that concurrent download is supported with your system's current firmware versions, using the information that is listed in Table 21.

Table 21. Firmware versions that support concurrent download (HP-UX)

DS4000 storage subsystem:	Controller firmware version:
DS4300 base and Turbo:	6.10.xx.xx. 6.12.xx.xx
DS4300 single-controller (SCU):	05.34.07.xx

Table 21. Firmware versions that support concurrent download (HP-UX) (continued)

DS4000 storage subsystem:	Controller firmware version:
DS4400/DS4500:	6.10.xx.xx, 6.12.xx.xx
DS4800	06.14.xx.xx
DS4100	06.10.xx.xx. 06.12.xx.xx
All other DS4000 storage subsystems:	05.30.xx.xx, 6.10.xx.xx, 6.12.xx.xx

Steps for upgrading the controller firmware

Complete the following steps to install a new firmware version on the controllers:

1. Download the correct version of the firmware from the Web site that is listed in "Check current versions of controller firmware and NVSRAM" on page 87. Place the file in a designated directory on the host system.

2. Check your system's current firmware versions to determine whether you need to stop I/O before upgrading the firmware, using the information that is listed in Table 21 on page 88.

   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.

3. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.

4. In the Subsystem Management window, click **Advanced —> Maintenance —> Download —> Controller Firmware**. The Download Firmware window displays.

5. Click **Browse** to select the firmware file that you previously downloaded to a directory on the host.

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**. The Confirm Download window displays.

8. Click **Yes** to start the download.

9. Close the Subsystem Management window and then reopen it to complete the firmware update.

10. Type the following command to restart the host:

```
# reboot
```

11. Verify that the firmware was successfully installed.

Steps for upgrading the NVSRAM

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

1. Download the correct NVSRAM version from the Web site that is listed in "Check current versions of controller firmware and NVSRAM" on page 87. Place the file in a designated directory on the host system.

2. Stop I/O.

3. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.

4. In the Subsystem Management window, click **Advanced —> Maintenance —> Download —> Controller NVSRAM**. The Download NVSRAM window displays.

5. Click **Browse** to select the NVSRAM file that you previously downloaded to a directory on the host.

6. Click **OK**. The Confirm Download window displays.

7. Click **Yes** to start the download.

8. Verify that the NVSRAM was successfully installed.

9. Start I/O.

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 24.

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 Topology section, highlight the storage subsystem and click **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 21 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 section, 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 Storage Subsystems 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 subsystems.
- 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 subsystems do not have minihubs.

- There are two or four HBAs (HP A6795A or HP A5158A) per DS4000 storage subsystem. 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 subsystem.
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 F, "Connecting HBAs in a FC switch environment," on page 223.

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 subsystem, 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 DS4000 controller firmware version 6.1.xx.xx or 5.xx.xx.xx installed, use method 1.
  - If you have DS4000 controller firmware 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 DS4000 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. 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 DS4000 controller has one logical drive access. For example, a subsystem that contains two DS4000 controllers has two logical drive accesses.

```
#SMdevices
```

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

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

3. Determine the preferred and alternate path for each logical drive by examining the output from the **SMdevices** command, as shown in the example in

Figure 11 on page 93. 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 on page 93, 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 on page 93, the alternate path is **/dev/rdsk/c172t0d0**:

```
/dev/rdsk/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 DS4000 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 on page 95. 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 DS4000 controller has one logical drive access. For example, a subsystem that contains two DS4000 controllers has two logical drive accesses.

```
#SMdevices
```

```
/dev/rdsk/c166t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0,
 Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>]
/dev/rdsk/c166t0d1 [Storage Subsystem DS4000, Logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>]
/dev/rdsk/c166t0d2 [Storage Subsystem DS4000, Logical Drive Finance, LUN 2,
Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>]
/dev/rdsk/c166t0d3 [Storage Subsystem DS4000, Logical Drive Purchasing, LUN 3,
Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>]
/dev/rdsk/c166t0d4 [Storage Subsystem DS4000, Logical Drive Development, LUN 4,
Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>]
/dev/rdsk/c166t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f56d00000001b00000000>]

/dev/rdsk/c172t0d0 [Storage Subsystem DS4000, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>]
/dev/rdsk/c172t0d1 [Storage Subsystem DS4000, logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>]
/dev/rdsk/c172t0d2 [Storage Subsystem DS4000, Logical Drive Finance, LUN 2,
Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>]
/dev/rdsk/c172t0d3 [Storage Subsystem DS4000, Logical Drive Purchasing, LUN 3,
Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>]
/dev/rdsk/c172t0d4 [Storage Subsystem DS4000, Logical Drive Development, LUN 4,
Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>]
/dev/rdsk/c172t3d7 [Storage Subsystem DS4000, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f5d6c0000002200000000>]
```

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
```

3. 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.

   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 subsystem. If you examine the WWNs for the
   logical drive access in Figure 12, you notice that they differ in only five digits,
   f56d0 and f5d6c.

   The devices in Figure 12 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:

   a. 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 <600a0b80000f56d00000001b00000000>]
```

   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>]
```

   b. 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>]
```

c. To keep a record for future reference, enter this path information for LUN 0 into a matrix (similar to the one in Table 22).

Table 22. 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 95 through step 3c 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.

Important: 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 22.

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/rdsk/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 96 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 host 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 subsystem.
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 102
- "Configuring storage subsystems" on page 108
- "Creating direct-attached and SAN-attached configurations" on page 114
- "Configuring the host bus adapter cards" on page 115
- "Installing a failover driver on Solaris" on page 125
- "Verifying external storage" on page 131
- "Redistributing logical drives in case of failure" on page 132

Hardware, firmware, and host system requirements

This section lists 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 subsystem. When the storage subsystem page opens, click the **Download** tab, and follow the links to the readmes.

DS4000 hardware and firmware requirements

Table 23 lists the DS4000 Fibre Channel storage subsystems that you can use with DS4000 Storage Manager.

Table 23. DS4000 storage subsystems that you can use with Solaris systems

Product name	Model	Controller firmware versions
IBM TotalStorage DS4800 Storage Subsytem	1815-82A/H, 1815-84A/H	06.14.xx.xx
IBM TotalStorage DS4100 Storage Subsytem	1724-100	06.12.xx.xx
IBM TotalStorage DS4100 Single-controller (SCU) Storage Subsytem	1724-1SX	05.42.xx.xx
IBM TotalStorage DS4500 Storage Subsytem	1742-90U	06.12.xx.xx
IBM TotalStorage DS4400 Storage Subsytem	1742-1RU	06.12.xx.xx
IBM TotalStorage DS4300 Standard and Turbo Storage Subsytem	1722-60U, 1722-60X	06.12.xx.xx

Table 23. DS4000 storage subsystems that you can use with Solaris systems (continued)

Product name	Model	Controller firmware versions
IBM TotalStorage DS4300 single-controller (SCU) Storage Subsytem	1722-6LU, 1722-6LX	5.34.xx.xx

NVSRAM: For the latest NVSRAM versions, see the most recent Storage Manager readme files for Solaris at the following Web site:
www-1.ibm.com/servers/storage/support/disk/
Click the link for your storage subsystem. When the storage subsystem page opens, click the **Download** tab, then click the Storage Manager link, then click the **Storage Mgr** tab.

Table 24 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 Enclosures, see *IBM TotalStorage DS4000 EXP100 Storage Expansion Enclosure Installation and User's Guide* (GC26-7601).

Table 24. Hardware you can use with Solaris systems

Hardware	Product Name	Model	Minimum firmware version
Storage Expansion Enclosure	IBM TotalStorage DS4000 EXP100 SATA Storage Expansion Enclosure	1710-10U or 10X	9554 **Note:** When upgrading the EXP100 to ESM firmware version 9554, you must use a storage subsystem that has DS4000 controller firmware **05.41.1x.xx** installed . You cannot upgrade to ESM firmware 9554 using a storage subsystem with different firmware installed.
	IBM TotalStorage DS4000 EXP710 Fibre Channel Storage Expansion Enclosure	1740-710	9627
	IBM TotalStorage DS4000 EXP700 Fibre Channel Storage Expansion Enclosure	1740-1RU or 1RX	9326 **Note:** You must upgrade the EXP700 ESM firmware to version 9326 before you upgrade the DS4000 controller firmware to version **06.1x.xx.xx**.
	IBM TotalStorage DS4000 EXP500 Fibre Channel Storage Expansion Enclosure	3560-1RU or 1RX	9166

Table 24. Hardware you can use with Solaris systems (continued)

Hardware	Product Name	Model	Minimum firmware version
	JNI FC Adapter	FCE-6410-N 33 Mhz	4.1.1.3
	JNI FC Adapter	FCE2-6412 66 Mhz 64 bit	
	JNI FC Adapter	FCE-6460-PCI 64 bit	5.3.1.1
	JNI FC Adapter	FCE-1473 S-bus 64 bit	
	JNI FC Adapter	FCC2-6562	
	JNI FC Adapter	FCX-6562	
	JNI FC Adapter	FCC\X2-6562	
	QLogic Adapter	4.08	2.00.01
	Emulex Adapter	LP9001	1.90.A4
	Emulex Adapter	LP9802	
	Emulex Adapter	LP10000	
Switches	IBM TotalStorage SAN switch	2109-S08, 2109-S16	2.6.2A
	IBM TotalStorage SAN switch	2109-F16	3.1.2A
	IBM TotalStorage SAN switch	2109-F8	
	IBM TotalStorage SAN switch	2109-M12	4.0.2C
	IBM TotalStorage SAN switch	2109-F32	4.2.0C
	McData switch	2032-064, 2032-140	6.01.00 - 18
	McData switch	2031-016	
	McData switch	2031-216, 2031-224	
	CNT switch	2042-001	4.3.0
	Cisco switch	2062-D01, 2062-D07, 2062-T07	1.3.4A

Note: If you have any questions about the latest DS4000 interoperability information, see the DS4000 interoperability matrix at the following Web site:

www-1.ibm.com/servers/storage/disk/ds4000/interop-matrix.html

Solaris system requirements

This section lists the minimum hardware and software requirements for a Solaris host in a DS4300, DS4400, DS4500, or DS4800 storage subsystem that is managed by DS4000 Storage Manager 9.1x.

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

Use the following software:

Operating system

Ensure that the Solaris host is running one of the following operating systems with the appropriate patches, including Sun's latest available recommended and security patches.

Solaris 8

Maintenance update 7

Solaris 9

Maintenance update 4, including the following patch:

- 113454-14 (or higher)

Solaris 10

VERITAS Cluster Server (optional) -- For high-availability clusters of Sparc S20 systems, install the VERITAS Cluster Server software package. See the VERITAS Cluster Server (VCS) documentation for the latest patches.

Note: The patches listed in this document can be superseded by more recent versions. Therefore, see your Solaris operating system documentation or contact your operating system supplier to ensure that you have the correct patches. In addition, review device driver readme files for other Solaris patches that you might need.

Installing DS4000 Storage Manager

The DS4000 Storage Manager software includes the following software packages:

SMruntime (required)

SMruntime is a Java compiler for the DS4000 Storage Manager software packages. SMruntime must be installed before the other packages are installed.

SMclient (optional)

SMclient is Storage Manager's graphical user interface (GUI), which you use to manage your storage subsystems. If you are not using your Solaris host as a storage management station (for example, if you are using a Windows management station with SMclient installed), then you do not need to install SMclient on your Solaris host.

SMagent (optional)

SMagent contains host-agent software, which you can use to manage your subsystems if you are using an in-band management configuration (managing the storage subsystems through the Fibre Channel I/O path to the host). If you are using an out-of-band management configuration, you do not need to install SMagent.

SMutil (required)

SMutil is a utility that you can use to register and map new logical drives to the operating system. SMutil is required on Solaris host systems.

You can install all of the DS4000 Storage Manager software packages automatically by using the Storage Manager installation wizard, or you can install each package manually. This section contains the following procedures:

- "Installing Storage Manager using the installation wizard"
- "Installing Storage Manager packages manually" on page 104 (original method)

Installing Storage Manager using the installation wizard

The DS4000 Storage Manager installation wizard can install all of the Storage Manager software packages automatically. If you prefer, you can install just the following **required** packages onto your Solaris host:

- SMruntime
- SMutil

To choose the packages that you want to install, select Custom Installation during the installation procedure, as described in "Steps for installing Storage Manager using the installation wizard."

Installation wizard requirements

To use the installation wizard, you must meet the following requirements:

- You are installing DS4000 Storage Manager 9.14
- Your Solaris host has a graphics adapter card installed.

If you are installing Storage Manager 9.10 or an earlier version on Solaris, or if your system does not have a graphics card installed, skip this section and install the stand-alone host software packages using the procedures described in "Installing Storage Manager packages manually" on page 104. All the packages are included with the installation CD.

Note: Do not install the software packages in the same server by using both the installation wizard and the manual installation method.

Steps for installing Storage Manager using the installation wizard

Complete the following steps to install the DS4000 Storage Manager software using the installation wizard. Adjust the steps as necessary for your specific installation.

1. Download the following file from the DS4000 Storage Manager CD, or from the DS4000 support Web site, to a directory on your system:

   SMIA-SOL-09.1x.xx.xx.bin

2. Open the file on your system. The Storage Manager Installation wizard's Introduction window opens.

3. Follow the instructions in each window of the wizard. When you select the installation type, you can choose one of the following options:

   - Typical (Full) Installation — Installs all Storage Manager software packages
   - Management Station — Installs SMruntime and SMclient
   - Host — Installs SMruntime, SMagent and SMutil
   - Custom — Allows you to select which packages you want to install

   Important: If you plan to install VERITAS DMP as a failover driver instead of RDAC, you must select **Custom Installation** as the installation

type. Then, when you select the components you want to install, make sure that RDAC is *not* selected.

4. Restart the system.

Result: The DS4000 Storage Manager software is installed on your system.

Installing Storage Manager packages manually

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 106

Note: If you are installing Storage Manager 9.12 or 9.14, do not install the software manually. Instead, use the installation procedure that is described in "Installing Storage Manager using the installation wizard" on page 103.

Installing the client software

Install the client software in the following order:

1. DS4000 SMruntime software (required)
2. DS4000 SMclient software (optional)

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, Solaris 9 or Solaris 10, and meets the minimum hardware and software requirements described in "Solaris system requirements" on page 101.
- 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 (required): 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 99 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 (optional):   If you are using your Solaris host as a storage management station, then you must install SMclient on the host. If you are using a different computer as your storage management station, install SMclient on that computer.

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, Solaris 9 or Solaris 10, and meets the minimum hardware and software requirements described in "Solaris system requirements" on page 101.
- 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 DS4000 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 subsystems 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 26 for more information.
- Versions of controller 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 Storage Manager readme file for Solaris 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 subsystem:

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

  Click the link for your storage subsystem. When the storage subsystem page opens, click the **Download** tab. Follow the links to the readmes.

  Note: If you do not have the latest versions of controller firmware and NVSRAM, download them and complete the procedure described in "Upgrading controller firmware and NVSRAM" on page 109. 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 Storage Subsystem**.
2. In the Add Storage Subsystem window, type the IP addresses of each controller in the storage subsystem, and click **Add**. The name of the storage subsystem displays in the Enterprise Management window.

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

Upgrading controller firmware and NVSRAM

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

Overview

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

- "Using concurrent download"
- "Using traditional or staged download"

Using concurrent download: Depending on your system's current firmware version, 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.

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 DS4000 controller firmware version 5.4x.xx.xx to version 6.1x.xx.xx. You can only use staged download if the DS4000 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 subsystems.

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 subsystem 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 subsystem that you plan to update.

Preparing for firmware or NVSRAM download

To prepare for download, review and complete all of the steps in this section before you start the download procedures.

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

To check the controller 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 view the current versions of controller firmware and NVSRAM.
3. Click **Close** to close the Storage Subsystem Profile window.
4. To find whether the versions of controller firmware and NVSRAM that are currently on your storage subsystem are the latest versions, see the Storage Manager readme file for Solaris at the following Web site:

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

   Click the link for your storage subsystem.

   When the storage subsystem page opens, click the **Download** tab.
5. If you find that you do not have the latest versions of the firmware and NVSRAM, download the correct versions from this Web site to a directory on your system, then follow the preparation and upgrade procedures that are described in this section.

Save the storage subsystem configuration: Saving the storage subsystem configuration to a script file is necessary in order to restore the configuration in the event of a catastrophic failure.

Attention: Do not use this option if the storage subsystem is undergoing an operation that changes any configuration settings. (If any logical drive listed in the Logical View of the Subsystem Management window has a clock icon, then a configuration change is occurring, and you must wait to save the configuration.)

You can also save the storage subsystem profile, by selecting **Storage Subsystem —> View Profile** in the Subsystem Management window, and clicking **Save As**. Then, in the Save Profile window, select the sections of the profile that you want to save, specify a name and location for the text file, and click **Save**.

Save the storage subsystem configuration by completing the following steps:

1. In the Subsystem Management window, click **Storage Subsystem —> Configuration —> Save**. The Save Configuration window displays.
2. In the Save Configuration window, select the elements of the configuration that you want to save, and click **Yes**. A save dialog displays.
3. In the save dialog, specify a name for the script file and select its location.

Notes:

- The .cfg extension is used by default for the script file if you do not specify a file extension.
- Do not store your DS4000 Storage Subsystem configuration script file in the same location as the logical drives that are defined on your DS4000 Storage Subsystem. If you do, and the logical drives fail during the subsequent migration, you will lose data.

4. Click **Save**. The Configuration Saved window displays.

5. Review the information in the Configuration Saved window, then click **OK** to return to the Subsystem Management window.

Determine whether you need to stop I/O: You can use concurrent download *only* if you are upgrading from one version of controller firmware to a later release of the same version.

Example: If you are upgrading from **06.10**.xx.xx to **6.12**.xx.xx, you can use concurrent download. However, if you are upgrading from **5.4x**.xx.xx to **06.1x**.xx.xx, you *cannot* use concurrent download.

In addition, before you upgrade controller firmware, ensure that concurrent download is supported with your system's current firmware versions, using the information that is listed in Table 25.

Table 25. Firmware versions that support concurrent download (Solaris)

DS4000 storage subsystem:	Controller firmware version:
DS4300 base and Turbo:	6.10.xx.xx. 6.12.xx.xx
DS4300 single-controller (SCU):	05.34.07.xx
DS4400/DS4500:	6.10.xx.xx, 6.12.xx.xx
DS4800	06.14.xx.xx
DS4100	06.10.xx.xx. 06.12.xx.xx
All other DS4000 storage subsystems:	05.30.xx.xx, 6.10.xx.xx, 6.12.xx.xx

Steps for upgrading the controller firmware

Complete the following steps to install a new firmware version on the controllers:

1. Download the correct version of the firmware from the Web site that is listed in "Check current versions of controller firmware and NVSRAM" on page 110. Place the file in a designated directory on the host system.

2. Check your system's current firmware versions to determine whether you need to stop I/O before upgrading the controller firmware, using the information that is listed in Table 25.

   Attention: If you do not have the latest firmware levels for concurrent download, you *must* stop I/O to the controllers before you upgrade the controller firmware.

3. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.

4. In the Subsystem Management window, click **Advanced —> Maintenance —> Download —> Controller Firmware**. The Download Firmware window displays.

5. Click **Browse** to select the firmware file that you previously downloaded to a directory on the host.

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**. The Confirm Download window displays.

8. Click **Yes** to start the download.

9. Close the Subsystem Management window and then reopen it to complete the firmware update.

10. Verify that the firmware was successfully installed.

Steps for upgrading the NVSRAM

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

1. Download the correct NVSRAM version from the Web site that is listed in "Check current versions of controller firmware and NVSRAM" on page 110. Place the file in a designated directory on the host system.

2. Stop I/O.

3. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.

4. In the Subsystem Management window, click **Advanced —> Maintenance —> Download —> Controller NVSRAM**. The Download NVSRAM window displays.

5. Click **Browse** to select the NVSRAM file that you previously downloaded to a directory on the host.

6. Click **OK**. The Confirm Download window displays.

7. Click **Yes** to start the download.

8. Verify that the NVSRAM was successfully installed.

9. Start I/O.

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 24.

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 Topology section of the Mappings window, highlight the name of the storage subsystem, and click **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 **Storage Subsystem —> Exit**.

Mapping LUNS to a storage partition

This section contains the following procedures:

- "Mapping LUNs to a new partition"
- "Mapping LUNs to an existing storage partition" on page 114

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

Mapping LUNs to a new partition

When mapping LUNs for the first time to a newly created partition, complete the following steps.

Before you begin: Read "DS4000 Storage Manager performance planning" on page 21 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 section, right-click the host on which you want to map LUNs, and select **Define Storage Partitioning**. The Define Storage Partitioning window opens.

3. In the Define Storage Partitioning window, select **Host**, then click **Next**.

4. Select the logical drive by name, on the right side of the window.

5. Accept the default LUN ID, or change it, then click **Add**.

6. Repeat step 5 for each LUN that you want to map to the partition.

Mapping LUNs to an existing storage partition

Before you begin: Read "DS4000 Storage Manager performance planning" on page 21 for information about planning the number of LUNs to map to each partition.

Complete the following steps to map LUNs to an existing storage partition:

1. Select the Mappings view of the Subsystem Management window.
2. In the Topology section, 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 Storage Subsystems 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 subsystems.

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 subsystems 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 subsystem.
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 F, "Connecting HBAs in a FC switch environment," on page 223.

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.

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 HBA settings and binding the HBAs" on page 116 (SAN-attached only)

QLogic
- "Installing the QLogic adapter driver package" on page 123
- "Modifying the QLogic HBA settings" on page 124

Installing the JNI adapter driver package

The procedures for installing the JNI adapter driver package and modifying the JNI HBA settings both refer to Table 26.

Table 26. JNI model numbers and configuration files

JNI HBA model numbers	Configuration files
FC64-1063-N	/kernel/drv/fcaw.conf
FCI-1063-N	/kernel/drv/fca-pci.conf
FCE2-1063, FCE2-1063, FCE-6410, FCE2-6412	/kernel/drv/jnic.conf
FCE-1473, FCE2-1473, FCE-6412 (2 Gb), FCE-6460	/kernel/drv/jnic146x.conf

Before you begin:
1. Ensure that the JNI HBAs have been installed and attached directly to the controllers or to a SAN fabric switch.
2. If attached to SAN fabric switch, ensure that the zones for the DS4000 storage subsystem have been created and enabled.

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

b. For more information about zoning and enabling zones, see Appendix F,
        "Connecting HBAs in a FC switch environment," on page 223.
3. In the Storage Manager readme file for Solaris, identify the correct driver
    packages for the installed HBAs. You can find the readme file at the following
    Web site:

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

    Click the link for your storage subsystem.

    When the storage subsystem page opens, click the **Download** tab.

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

1. Download the most current HBA driver package, as listed in the Storage
    Manager readme file for Solaris, from the following Web site:

    www.jni.com/Drivers

2. 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
```

3. 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
```

4. Type the following command to reboot the Solaris host:

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

Modifying the HBA settings and binding the HBAs

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 DS4000 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, CNT, and IBM SAN 1 GB
    switches" on page 117
*   "Modifying the JNI HBA settings for a IBM SAN 2 Gb switch" on page 120

Attention: If you have a direct-attached configuration, skip this section and proceed
to "Installing the RDAC failover driver" on page 125 or "Installing the DMP failover
driver" on page 127. You do not need to modify JNI HBA settings if you have a
direct-attached configuration.

Modifying the JNI HBA settings for McData, Cisco, CNT, and IBM SAN 1 GB switches

Before you begin: Table 26 on page 115 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 201.

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 that is listed in Appendix B, "JNI and QLogic host bus adapter settings," on page 201.

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, as in the following example:

```
Jul 18 14:33:50 sunburn jnic146x: [ID 550325 kern.notice] jnic146x0:
Hba: JNI,FCR     Model: FCX2-6562
Jul 18 14:33:50 sunburn jnic146x: [ID 277434 kern.notice] jnic146x0:
FCode: Version 3.8.9 [BAFA]
Jul 18 14:33:50 sunburn jnic146x: [ID 188160 kern.notice] jnic146x0:
IRQ: 1  Channel: 1
Jul 18 14:33:50 sunburn jnic146x: [ID 428106 kern.notice] jnic146x0:
SubVend: 1242  Sub: 6562  Rev: b100  Vend: 124 2  Dev: 1560
Jul 18 14:33:50 sunburn jnic146x: [ID 904290 kern.notice] jnic146x0:
WWNN: 1000000173008A2E  WWPN: 2000000173008A2E
Jul 18 14:33:51 sunburn jnic146x: [ID 695386 kern.notice] jnic146x0:
Auto-negotiating link speed
Jul 18 14:33:51 sunburn jnic146x: [ID 873028 kern.notice] jnic146x0:
Configured as Private Loop port
Jul 18 14:33:52 sunburn jnic146x: [ID 245684 kern.notice] jnic146x0:
JNIC v5.2.3 (03042900)

Jul 18 14:33:50 sunburn jnic146x: [ID 550325 kern.notice] jnic146x1:
Hba: JNI,FCR     Model: FCX2-6562
Jul 18 14:33:50 sunburn jnic146x: [ID 277434 kern.notice] jnic146x1:
FCode: Version 3.8.9 [BAFA]
Jul 18 14:33:50 sunburn jnic146x: [ID 188160 kern.notice] jnic146x1:
IRQ: 1  Channel: 1
Jul 18 14:33:50 sunburn jnic146x: [ID 428106 kern.notice] jnic146x1:
SubVend: 1242  Sub: 6562  Rev: b100  Vend: 124 2  Dev: 1560
Jul 18 14:33:50 sunburn jnic146x: [ID 904290 kern.notice] jnic146x1:
WWNN: 1000000173009A2E  WWPN: 2000000173009A2E
Jul 18 14:33:51 sunburn jnic146x: [ID 695386 kern.notice] jnic146x1:
Auto-negotiating link speed
Jul 18 14:33:51 sunburn jnic146x: [ID 873028 kern.notice] jnic146x1:
Configured as Private Loop port
Jul 18 14:33:52 sunburn jnic146x: [ID 245684 kern.notice] jnic146x1:
JNIC v5.2.3 (03042900)
```

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. Record the following information about the installed HBAs:
   - The JNI driver instances
   - The WWPNs

   For example, for the output listed above, you would record the following information:

JNI driver instance	WWPN
jnic146x0	2000000173008A2E
jnic146x1	2000000173009A2E

After you finish gathering the binding information, continue to the next procedure, "Binding the HBAs to storage controllers."

Binding the HBAs to storage controllers: The examples in this procedure assume that you are binding two JNIC FCE-6460 HBAs to the DS4000 controllers. Adjust the steps for your specific configuration.

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. In the file, search for a section called `Configuration parameters for target to FC device mappings`, as shown in the following example:

```
##############################################################################
#
# Configuration parameters for target to FC device mapping:
#
#     targetX_hba,
#     targetX_lunY_hba,
#     targetX_wwnn,
#     targetX_wwpn,
#     targetX_port,
#
# Type: string
# Default:
#
# Example usage: Replace the 'x' characters with actual WWNN, WWPN, or port ID.
#
# target0_hba      = "jnic146x0";
# target0_lun0_hba = "jnic146x0";
# target0_wwnn     = "xxxxxxxxxxxxxxxx";
# target0_wwpn     = "xxxxxxxxxxxxxxxx";
# target0_port     = "xxxxxx";
# jnic146x0-target0_lun0_hba="jnic146x0";
# jnic146x0-target0_wwnn="xxxxxxxxxxxxxxxx";



# Example usage: Replace the 'x' characters with actual WWNN, WWPN, or port ID.
#
# target0_hba      = "jnic146x0";
# target0_lun0_hba = "jnic146x0";
# target0_wwnn     = "xxxxxxxxxxxxxxxx";
# target0_wwpn     = "xxxxxxxxxxxxxxxx";
# target0_port     = "xxxxxx";
# jnic146x0-target0_lun0_hba="jnic146x0";
# jnic146x0-target0_wwnn="xxxxxxxxxxxxxxxx";
```

3. Using the information that you recorded in the previous procedure, "Gathering the binding information" on page 117, add the following parameters to the file for each HBA:

   • Target*x*_hba = JNI driver instance
   • Target*x*_wwpn = Controller WWPN

For example:

```
Target0_hba = jnic146x0
Target0_wwpn = 2000000173008A2E

Target1_hba = jnic146x1
Target1_wwpn = 2000000173009A2E
```

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

```
#:wq
```

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

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

Modifying the JNI HBA settings for a IBM SAN 2 Gb switch

If you are using a IBM SAN 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, CNT, and IBM SAN 1 GB switches" on page 117 for the correct modification procedures.

Complete the following procedures to modify the JNI HBA settings.

Gathering the binding information:

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:

```
# /SunOS
```

3. Look for the HBAs in the /var/adm/messages file. For example, you might see the following information in the /var/adm/messages file:

```
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.
```

4. Record the following information about the installed HBAs:
   - The JNI driver instances
   - The WWPNs

   For example, for the output listed above, you would record the following information:

JNI driver instance	WWPN
jnic146x0	200300a0ba0f478F
jnic146x3	200200a0ba0f478F

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`:

   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 the JNI configuration file: 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 201 and modify settings accordingly.

1. Type the following command:

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

2. In the file, 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, 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: IBM SAN 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. In the jnic146x.conf file, search for a section called `Configuration parameters for target to FC device mappings`, as shown in the following example:

```
#########################################################################
#
# Configuration parameters for target to FC device mapping:
#
#     targetX_hba,
#     targetX_lunY_hba,
#     targetX_wwnn,
#     targetX_wwpn,
#     targetX_port,
#
# Type: string
# Default:
#
# Example usage: Replace the 'x' characters with actual WWNN, WWPN, or port ID.
#
# target0_hba       = "jnic146x0";
# target0_lun0_hba = "jnic146x0";
# target0_wwnn      = "xxxxxxxxxxxxxxxx";
# target0_wwpn      = "xxxxxxxxxxxxxxxx";
# target0_port      = "xxxxxx";
# jnic146x0-target0_lun0_hba="jnic146x0";
# jnic146x0-target0_wwnn="xxxxxxxxxxxxxxxx";



# Example usage: Replace the 'x' characters with actual WWNN, WWPN, or port ID.
#
# target0_hba       = "jnic146x0";
# target0_lun0_hba = "jnic146x0";
# target0_wwnn      = "xxxxxxxxxxxxxxxx";
# target0_wwpn      = "xxxxxxxxxxxxxxxx";
# target0_port      = "xxxxxx";
# jnic146x0-target0_lun0_hba="jnic146x0";
# jnic146x0-target0_wwnn="xxxxxxxxxxxxxxxx";
```

2. Using the information that you recorded in the previous procedure, "Gathering the binding information" on page 120, add the following parameters to the file for each HBA:

   - Target*x*_hba = JNI driver instance
   - Target*x*_wwpn = Controller WWPN

   For example:

```
Target0_hba = jnic146x0
Target0_wwpn = 200300a0ba0f478F

Target1_hba = jnic146x3
Target1_wwpn = 200200a0ba0f478F
```

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

```
#:wq
```

4. 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 125 or "Installing the DMP failover driver" on page 127 for more information.

Installing the QLogic adapter driver package

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

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 F, "Connecting HBAs in a FC switch environment," on page 223.

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 201.

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 127

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 115.

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 Storage Manager readme file for Solaris 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 subsystem. When the storage subsystem page opens, click the **Download** tab.

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 HBA settings and binding the HBAs" on page 116.

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.

Important This section refers to the following Web site for some installation procedures:

www.veritas.com

System requirements

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

- Solaris operating system
- VERITAS Volume Manager 3.5 or 4.1
- 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 F, "Connecting HBAs in a FC switch environment," on page 223.

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 130
3. "Installing the SMibmasl software package" on page 130
4. "Configuring VERITAS Volume Manager" on page 130
5. "Starting the VERITAS Enterprise Administrator" on page 131

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 fil, 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 112.

    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

Before you begin to install VERITAS Volume Manager packages for DMP, 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.

This document does not describe how to install the VERITAS product. For VERITAS Volume Manager installation procedures, see the VERITAS product documentation for Solaris at the following Web site:

www.veritas.com

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

This document does not describe how to configure the VERITAS product. For VERITAS Volume Manager configuration procedures, see the VERITAS product documentation for Solaris at the following Web site:

www.veritas.com

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 107.

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 expansion drawers 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 subsystem.
2. Using the Subsystem Management window, redistribute logical drives to their preferred paths by clicking **Advanced —> Recovery —> Redistribute Logical Drives**.

Chapter 5. Installing and configuring storage management software on POWER-based Linux host systems

You can use DS4000 Storage Manager with the following Linux distributions for POWER-based host operating systems:

- SUSE Enterprise Linux Server 9 (SLES 9)
- Red Hat Enterprise Linux 3 (RHEL 3)

This chapter contains the following topics, which describe how to install and use DS4000 Storage Manager on SLES 9 and RHEL 3 hosts:

- "Hardware, firmware and system requirements"
- "Restrictions" on page 138
- "Downloading the HBA driver and utility packages" on page 139
- "Installing DS4000 Storage Manager and utilities" on page 143
- "Configuring storage subsystems" on page 149
- "Installing the RDAC Multipath Proxy driver" on page 161
- "Performing initial device identification" on page 163
- "Performing common maintenance tasks" on page 169

For additional 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, firmware and system requirements

This section lists the following requirements:

- Hardware and firmware requirements for the DS4000 configuration
- Hardware and software requirements for your POWER-based SLES 9 or RHEL 3 host system

Important: For the latest installation requirements and user information, see the Storage Manager readme file for Linux on POWER that is located in the 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 DS4000 Mid-range disk system. When the storage subsystem page opens, click the **Download** tab, and follow the links to the Storage Manager readmes.

DS4000 requirements

Table 27 on page 134 lists the DS4000 Storage Subsystems that you can use with POWER-based SLES 9 or RHEL 3 hosts.

Table 27. DS4000 storage subsystems that you can use with POWER-based Linux host systems

Product name	Model	Controller firmware versions
IBM TotalStorage DS4800 Storage Subsystem	1815-82A/H, 1815-84A/H	06.14.xx.xx
IBM TotalStorage DS4100 Storage Subsystem	1724-100	06.12.xx.xx
IBM TotalStorage DS4100 Single-controller (SCU) Storage Subsystem	1724-1SX	05.42.xx.xx
IBM TotalStorage DS4500 Storage Subsystem	1742-90U	06.12.xx.xx
IBM TotalStorage DS4400 Storage Subsystem	1742-1RU	06.12.xx.xx
IBM TotalStorage DS4300 Storage Subsystem	1722-60U, 1722-60X	06.12.xx.xx
IBM TotalStorage DS4300 single-controller (SCU) Storage Subsystem	1722-6LU, 1722-6LX	5.34.xx.xx

NVSRAM: For the latest NVSRAM versions, see the most recent Storage Manager readme files for Linux on POWER at the following Web site:
www-1.ibm.com/servers/storage/support/disk/
Click the link for your storage subsystem. When the storage subsystem page opens, click the **Download** tab, then click the Storage Manager link, then click the **Storage Mgr** tab.

Table 28 on page 135 lists the DS4000 EXP storage expansion enclosures and other hardware that are supported with POWER-based SLES 9 or RHEL 3 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 28. Hardware that is supported with POWER-based Linux hosts

Hardware	Product Name	Model	Minimum firmware version
Storage Expansion Enclosures	IBM TotalStorage DS4000 EXP100 SATA Storage Expansion Enclosure	1710–10U or 10X	9554 **Note:** When upgrading the EXP100 to ESM firmware version 9554, you must use a storage subsystem that has DS4000 controller firmware **05.41.1x.xx** installed . You cannot upgrade to ESM firmware 9554 using a storage subsystem with different firmware installed.
	IBM TotalStorage DS4000 EXP710 Fibre Channel Storage Expansion Enclosure	1740-710	9627
	IBM TotalStorage DS4000 EXP700 Fibre Channel Storage Expansion Enclosure	1740-1RU or 1RX	9326 **Note:** You must upgrade the EXP700 ESM firmware to version 9326 before you upgrade the DS4000 controller firmware to version **06.1x.xx.xx**.
	IBM TotalStorage DS4000 EXP500 Fibre Channel Storage Expansion Enclosure	3560-1RU or 1RX	9166
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	1.90A4

Table 28. Hardware that is supported with POWER-based Linux hosts (continued)

Hardware	Product Name	Model	Minimum firmware version
Switches	IBM TotalStorage SAN switch	2109-S08 2109-S16	2.6.2A
	IBM TotalStorage SAN switch	2109-F16	3.1.2A
	IBM TotalStorage SAN switch	2109-M12	4.0.2C
	IBM TotalStorage SAN switch	2109-F32	4.2.0C
	IBM TotalStorage SAN 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
	CNT switch	2042-001	4.3.0
	Cisco switch	2062-D01, 2062-D07, 2062-T07	1.3.4A

Check the latest Storage Manager readme file for Linux on POWER for the most up-to-date firmware versions.

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 subsystem. When the storage subsystem page opens, click the **Download** tab, and follow the links to the Storage Manager readmes.

Note: If you have any questions about the latest DS4000 interoperability information, see the DS4000 interoperability matrix at the following Web site:

www-1.ibm.com/servers/storage/disk/ds4000/interop-matrix.html

SUSE Enterprise Linux Server (SLES 9) system requirements

This section lists the minimum hardware and software requirements for a POWER-based host running SLES 9 on a DS4800, DS4300 Turbo, DS4400, or DS4500 storage subsystem that is managed by DS4000 Storage Manager.

Host hardware requirements

The following hardware is required at minimum:

- **Host:** IBM @server pSeries or POWER server with 64-bit PowerPC architecture™
- **Management station for DS4000 Storage Manager 9.14**

  You can use the SLES 9 host as the management station for DS4000 Storage Manager 9.14.

- **Management station for DS4000 Storage Manager 9.12 and earlier**

  DS4000 Storage Manager 9.12 (and earlier) client software and DS4000 controller firmware 06.12.xx.xx are not supported on SLES 9 hosts, so you must install the Storage Manager software on a separate 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, Solaris, or Red Hat 3 and Red Hat 4 management stations, see the minimum requirements listed in the appropriate chapters in this document.

Software requirements

The following software distributions must be at the specified versions or later:

SUSE LINUX Enterprise Server 9
- **Version:** SP1
- **Kernel:** 2.6.5-7.139-pseries64

Emulex HBA driver and Application Kit
- lpfcdd v2.10g driver (included with SLES 9 installation CD)
- lpfcdd v2.10h driver
- Emulex v2.10g Application Kit

IBM Storage Manager RDAC MPP driver (firmware 06.14.xx.xx only)
rdac-LINUX-09.01.B5.02-source.tar.gz

IBM Storage Manager software packages (firmware 06.14.xx.xx only)
SMIA-LINUXPPC-09.14.A5.05.bin

The following individual software packages are included in the SMIA package. You can install them with the SMIA package, or separately:
- SMruntime-LINUX-09.14.A5.00-1.ppc64.rpm
- SMclient-LINUX-09.14.G5.05-1.noarch.rpm
- SMutil-LINUX-09.14.A5.00-1.ppc64.rpm
- SMagent-LINUX-09.14.A5.00-1.ppc64.rpm

Red Hat Enterprise Linux system requirements

This section lists the minimum hardware and software requirements for a POWER-based host running RHEL 3 (update 4) on a DS4300 Turbo, DS4400, or DS4500 storage subsystem that is managed by DS4000 Storage Manager.

Host hardware requirements

The following hardware is required at minimum:

Host: IBM @server pSeries or POWER server with 64-bit PowerPC architecture.

Software requirements

The following software must be at the specified versions, or later. Check the Storage Manager readme file for Linux on POWER at the following Web site for the most up-to-date version levels.

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

Click the link for your storage subsystem. When the storage subsystem page opens, click the **Download** tab, and follow the links to the readmes.

Red Hat Enterprise Linux
- Version 3, Update 4
- Kernel 2.4.21-27.EL

Emulex tool kit and HBA driver
- lpfc v7.1.14 (included with kernel)
- ioctl kit (Helper Module: lpfc_2.4_ioctl.module_kit — 1.0.13–1.tar.gz)
- Emulex v7.1.14 Application Kit (ElxLinuxApps — 2.0a20-7.1.1.14-1.tar)

IBM Storage Manager RDAC MPP driver (for failover capability)
rdac-LINUX-09.00.A5.09-source.tar.gz

IBM Storage Manager software packages
SMIA-LINUXPPC-09.14.A5.05.bin

The following individual software packages are included in the SMIA package. You can install them with the SMIA package, or separately:
- SMruntime-LINUX-09.14.A5.00-1.ppc64.rpm
- SMclient-LINUX-09.14.G5.05-1.noarch.rpm
- SMutil-LINUX-09.14.A5.00-1.ppc64.rpm
- SMagent-LINUX-09.14.A5.00-1.ppc64.rpm

Restrictions

Before you begin working with DS4000 Storage Manager, read the following list of restrictions. The restrictions apply to both SLES 9 and RHEL 3, except as otherwise noted:

SAN and connectivity restrictions
- Linux on POWER-based hosts can support the following number of host bus adapters (HBAs):
  - **Single-path:** Single-path configurations can support up to two HBAs 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 and a separate partition.
  - **Dual-path:** Dual-path configurations can support two HBAs per partition, and one partition per host, per DS4000 storage subsystem.
- Direct-attached single-path configurations are restricted to single-initiator configurations only. You can only make one connection to each minihub on a DS4000 storage subsystem.

  Note: DS4300 Turbo storage subsystems 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.
- Clustering is not currently supported.

  Note: For more information about zoning and enabling zones, see Appendix F, "Connecting HBAs in a FC switch environment," on page 223.

Linux RDAC restrictions
- Auto volume transfer/auto-disk transfer (AVT/ADT) is managed by the Linux RDAC driver. To avoid contention, you must disable AVT/ADT by selecting the `LNXCL` host type instead of the `Linux` host type when you define your host groups. (The `Linux` host type automatically enables AVT/ADT; the `LNXCL` host type does not.)

Partitioning restrictions

- The maximum number of partitions per POWER-based Linux host, per DS4000 storage subsystem, is as follows:
  - **Single-path:** Four partitions
  - **Dual-path:** One partition
- Map the host type as follows:
  - **Single-path:** All logical drives in single-path configurations must be mapped as host type `Linux` (AVT is enabled in this host type).
  - **Dual-path:** All logical drives in dual-path configurations must be mapped as host type `LNXCL` (AVT is disabled in this host type).
- The Linux SCSI layer does not support skipped (sparse) LUNs. If the mapped LUNs are not contiguous, the Linux kernel does not scan any LUNs that occur after the skipped LUN, so the remaining LUNs are not available to the host server. Therefore, always map the LUNs using consecutive numbers, starting with LUN 0. For example, map LUNs to numbers 0; 1; 2; 3; 4; 5; and so on, without skipping any numbers.
- On each partition, you must map a LUN 0.
- On each controller, you must configure at least one partition with an ID between 0 and 31 that is not a UTM or access logical drive.

For more information about Storage Partitioning, see "Storage Partitioning premium feature" on page 24.

Interoperability restrictions

- **Important:** Do not attach DS4000 devices before the first time you boot the host.
- Dynamic volume expansion (DVE) is not currently supported on RHEL 3. (DVE is supported on SLES 9.)

Downloading the HBA driver and utility packages

This section contains the following procedures:

- "Downloading the SLES 9 pSeries HBA driver and utility packages"
- "Downloading the RHEL 3 HBA driver utility packages" on page 141

Downloading the SLES 9 pSeries HBA driver and utility packages

Download the following packages from the Emulex Web site:

www.emulex.com/ts/docoem/framibm.htm

	Compressed package filename	Uncompressed filename	Description
Application package	lpfcapp-2.10g-1.ppc64.tgz	lputil	Monitors HBAs and facilitates downloading of HBA firmware
Driver package	lpfcdriver-2.10h.tgz	lpfcdd driver	Compiles changes in the configuration file (/etc/lpfc.conf)
Release notes	(Not compressed)	linux_2.0h.pdf	Instructions for upgrading to the 2.10h device driver

	Compressed package filename	Uncompressed filename	Description

Note: You can use the **lputil** utility to check current HBA firmware levels or to download firmware.

For each compressed package, type the following command to uncompress the files:

```
# tar -zxvf package_name.tgz
```

Notes:

1. Refer to the release notes for lpfc driver package installation procedures.
2. If you are not installing RDAC at this time, see the release notes for information about making a new initrd image.

To install the lpfc driver, complete whichever of the following procedures is appropriate for your configuration:

- "Installing the lpfc driver package on a host without RDAC installed"
- "Installing the lpfc driver package on a host with RDAC installed" on page 141

Installing the lpfc driver package on a host without RDAC installed

Use this procedure if you are installing the lpfc driver on a host that does not have RDAC already installed. If RDAC is installed, see "Installing the lpfc driver package on a host with RDAC installed" on page 141.

Complete the following steps to install the lpfc driver.

1. Create a directory for the lpfc driver.
2. Download the **lpfcdriver-2.10h.tgz** package file to the directory you just created.
3. Uncompress the file by typing the following command:

```
# tar -zxvf package_name.tgz
```

4. In the same directory, install the module by typing the following commands:

```
# sh Install.sh
# make
# make modules_install
```

5. If a previous version of the lpfcdd module is loaded, unload the that lpfcdd module by typing the following command. (If a previous version of the lpfcdd module is not on the host, go to step 6.)

```
# rmmod lpfcdd
```

6. Load the new module by typing the following command:

```
# insmod lpfcdd.ko
```

7. Edit the /etc/yaboot.conf directory, as appropriate.
8. Reboot.

Installing the lpfc driver package on a host with RDAC installed

If the host already has RDAC installed, you must uninstall RDAC before installing the lpfc driver.

Note: If the RDAC driver is not installed, and you are not installing the RDAC driver at this time, see the release notes for information about making a new initrd image.

Complete the following steps to uninstall RDAC and install the lpfc driver:

1. Change to the Linux RDAC directory by typing the following command:

```
# cd ~linuxrdac dir
```

2. Uninstall RDAC by typing the following command:

```
# make uninstall
```

3. Create a directory for the lpfc driver.
4. Download the **lpfcdriver-2.10h.tgz** package file to the directory you just created.
5. Uncompress the file by typing the following command:

```
# tar -zxvf package_name.tgz
```

6. In the same directory, install the module by typing the following commands:

```
# sh Install.sh
# make
# make modules_install
```

7. If a previous version of the lpfcdd module is loaded, unload the that lpfcdd module by typing the following command. (If a previous version of the lpfcdd module is not on the host, go to step 8.)

```
# rmmod lpfcdd
```

8. Load the new module by typing the following command:

```
# insmod lpfcdd.ko
```

9. Reinstall the RDAC driver by typing the following commands:

```
# cd ~linuxrdac dir
# make clean
# make
# make install
```

10. Edit the /etc/yaboot.conf directory, as appropriate.
11. Reboot.

Downloading the RHEL 3 HBA driver utility packages

The Emulex Helper Module and Application Kit are not included with the RHEL 3 kernel. You need to download them, as described in this section.

Note: The Emulex HBA driver is included with the RHEL 3 kernel.

Complete the following steps to download and install the Emulex HBA Kits.

1. Download the following packages from the Emulex Web site:

   www.emulex.com/ts/docoem/framibm.htm

Description	Compressed package filename
Emulex Helper Module	lpfc_2.4_ioctl_module_kit-1.0.13-1.tar.gz
Emulex Application Kit	ElxLinuxApps-2.0a20-7.1.14-1.tar

2. Type the following command to uncompress the Emulex Helper Module:

```
# tar -zxvf lpfc_2.4_ioctl_module_kit-1.0.13-1.tar.gz
```

   Result: The uncompressed package file is put into a new directory.

3. Type the following commands to install the module:

```
# cd lpfc_2.4_ioctl_module_kit-1.0.13-1
# ioctl-install
```

4. Type the following command to uncompress the Emulex Application Kit:

```
# tar -xvf ElxLinuxApps-2.0a20-7.1.14-1.tar
```

   Result: The package file is uncompressed.

5. Type the following command to install the kit:

```
# ./install
```

   Important: HBAnyware starts automatically when you install the Emulex
   Application Kit. Because HBAnyware impedes the removal of the
   lpfcdfc file from the modules stack when using **insmod**, you must
   stop HBAnyware by using the **stop hbanyware** or **uninstall**
   commands at these times. For more information, see the readme.txt
   file that accompanies the Application Kit files.

Result: The following files and directories are installed:

- enterprise_kitfiles/
- enterprise_kitfiles/hbaapi_2.0.f.tgz (loads HBAnyware)
- enterprise_kitfiles/lputil_16a10.tgz
- enterprise_kitfiles/EmlxApps300a15.tgz
- enterprise_kitfiles/lpfcutil_2.4_lib_kit-1.0.13-1.tgz
- install
- uninstall
- README.txt
- README_SSC.txt

Note: Use the **lputil** utility to check current HBA firmware levels or to download
firmware.

Installing DS4000 Storage Manager and utilities

With Storage Manager 9.14, the Storage Manager client software and utilities are available for POWER-based RHEL 3 and SLES 9 hosts.

Exception: On storage subsystems running controller firmware 9.12.xx.xx, you *cannot* use your POWER-based SLES 9 host as a storage management station. To manage these DS4000 storage subsystems with your POWER-based SLES 9 host, you must install the Storage Manager client software (SMclient) on an AIX 5L, Solaris, HP-UX, Windows, i386 processor-based Linux, or POWER-based RHEL 3 server, and use that server as your storage management station.

- For information about 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 27
  - Chapter 3, "Installing and configuring storage management software on HP-UX host systems," on page 77
  - Chapter 4, "Installing and configuring storage management software on Solaris host systems," on page 99
- For information about 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 Intel-based Operating System Environments* (GC26-7649).
- For information about installing DS4000 Storage Manager on a POWER-based RHEL 3 management station, see the next section, "Prerequisites for installing Storage Manager on a RHEL 3 host" on page 144.

Prerequisites for installing Storage Manager on a SLES 9 host

After you install SLES 9 on your host, and before you begin to install the Storage Manager software and configure your subsystems, complete the following tasks:

1. "Downloading the SLES 9 pSeries HBA driver and utility packages" on page 139
2. "Checking the /etc/sysconfig/kernel file"
3. "Rebuilding the ramdisk" (required only if changes are made to the /etc/sysconfig/kernel file)

Important: When you boot the host the first time, you must not have any DS4000 devices attached.

Checking the /etc/sysconfig/kernel file

Ensure that the local hard disk driver is listed before the lpfcdd driver in the /etc/sysconfig/kernel file, as in the following example:

```
INITRD_MODULES="sym53c8xx lpfcdd"
```

where *sym53c8xx* is the local hard disk driver. If the listings are not in that order, edit the file to list the lpfcdd driver last.

Note: If you edit this file, you must rebuild the ramdisk, as described in "Rebuilding the ramdisk."

Rebuilding the ramdisk

If you edited the /etc/sysconfig/kernel file, you must complete the following steps:

1. Type the following commands:

```
# cd /boot
# mkinitrd -v maxscsi.initrd-2.4.21-27.EL.img 2.4.21-27.EL
```

where *maxscsi* is an arbitrary name and *2.4.21-27.EL* is the current kernel version.

2. Edit the /etc/yaboot.conf file by adding the following lines:

```
image=/vmlinux-2.4.21-27.EL
    label=max-linux
    read-only
    initrd=/maxscsi.initrd-2.4.21-27.EL.img
    append="root=LABEL=/"
```

where *maxscsi* is the name you used in step 1 and *2.4.21-27.EL* is the current kernel version.

Important: If you do not complete this step, the outdated initrd space will be installed when you reboot.

3. Reboot.

Prerequisites for installing Storage Manager on a RHEL 3 host

After you install RHEL 3 on your host, and before you begin to install the Storage Manager software and configure your subsystems, you must complete the following tasks:

- "Modifying the configuration file"
- "Rebuilding the ramdisk" on page 145
- "Downloading the RHEL 3 HBA driver utility packages" on page 141

Important: When you boot the host the first time, you must not have any DS4000 devices attached.

Modifying the configuration file

Complete the following steps to modify the configuration file (/etc/modules.conf):

1. On your RHEL 3 host, type the following command to view the configuration file:

```
# view /etc/modules.conf
```

2. In the file, verify that the lpfc device driver is listed *after* all the other scsi device drivers. If it is not, change the listing order and the scsi_hostadapter number so that the lpfc device driver is listed last, as in the following example:

```
# view /etc/modules.conf
alias eth0 e100
alias eth1 e100
alias eth2 e100
alias scsi_hostadapter sym53c8xx
alias scsi_hostadapter1 lpfc
```

Important: If you do not complete this step, and the lpfc module is listed before the internal disk module, the DS4000 devices will be discovered first. Name slippage of the boot disk will occur.

3. Add the following lines to the configuration file:

```
options scsi_mod max_scsi_luns=255
options lpfc lpfc_nodev_tmo=60 lpfc_lun_queue_depth=10 fc_inq_pqb_filter=0x01
```

   Explanation of new lines:
   - `scsi_mod max_scsi_luns=255` — This line is required so that the host can read more than one device on each storage subsystem
   - `lpfc lpfc_nodev_tmo=60` — The setting is not adjustable. 60 is the required setting.
   - `lpfc_lun_queue_depth=xx` The setting is adjustable. The default is 30.
   - `lpfc_inq_pqb_filter=0x01` Use this setting if needed. The filter parameter fixes a known SCSI midlayer problem in some distributions. Use this setting if you find duplicate LUN 0's in /proc/scsi/scsi, per attached DS4000.

After you finish updating the configuration file, continue to the next procedure, "Rebuilding the ramdisk."

Rebuilding the ramdisk

After you modify the configuration file, complete the following steps to rebuild the ramdisk:

1. Update the modules.dep file by typing the following command:

```
# depmod -a
```

2. Type the following commands:

```
# cd /boot
# mkinitrd -v maxscsi.initrd-2.4.21-27.EL.img 2.4.21-27.EL
```

   where *maxscsi* is an arbitrary name and *2.4.21-27.EL* is the current kernel version.

3. Edit the /etc/yaboot.conf file by adding the following lines:

```
image=/vmlinux-2.4.21-27.EL
    label=max-linux
    read-only
    initrd=/maxscsi.initrd-2.4.21-27.EL.img
    append="root=LABEL=/"
```

   where *maxscsi* is the name you used in step 2 and *2.4.21-27.EL* is the current kernel version.

   Important: If you do not complete this step, the outdated initrd space will be installed when you reboot.

4. Reboot.

Installing Storage Manager on a POWER-based Linux host

Use the procedures in this section to install DS4000 Storage Manager 9.14 on your SLES 9 or RHEL 3 host.

Note: DS4000 Storage Manager versions 9.12 and 9.10 are not supported on SLES 9 hosts.

You can install all of the DS4000 Storage Manager 9.14 software packages automatically by using the Storage Manager installation wizard, or you can install each package manually. This section contains the following procedures:

- "Installing Storage Manager using the installation wizard"
- "Installing Storage Manager packages manually" on page 147 (original method)

Installing Storage Manager using the installation wizard

The DS4000 Storage Manager installation wizard is a Java-based interactive method of choosing which packages to automatically install on your host system. The installation wizard installs the following software packages on the SLES 9 or RHEL 3 host:

- SMruntime
- SMclient
- SMagent
- SMutil

Note: If you are installing Storage Manager 9.10 or 9.12, if your system does not meet the requirements for using the wizard, or if for any reason you do not want to use the wizard to install the software, skip this section and install the stand-alone host software packages using the procedures described in "Installing Storage Manager packages manually" on page 147. All the packages are included with the installation CD.

Requirements:

- DS4000 Storage Manager 9.14
- Graphics adapter card or a remote graphical console program (such as VNC or Exceed)
- Graphical Interface X Configuration (such as XFree86) installed and configured on the host

Opening a remote Xwindows or VNC connection: The Storage Manager 9.14 SMIA package must be installed in a graphics environment, as indicated in the listed requirements. If the local console does not have a graphics card, you can install the package remotely through VNC or Xwindows connections, as follows:

Xwindows connection

Complete the following steps to open a remote Xwindow connection:

1. Ensure that an Xwindows program is installed and running on your server.
2. Open a remote Xwindow from your PC using Exceed by typing the following command line in the Exceed Client Startup window with appropriate other selections:

```
 xterm -display PC-IP-ADDR:0 -T remote-IPADDR
-r -ls +sb -sl 2000 -n remote-IPADDR -geom 80x24+0+0
```

VNC connection

Complete the following steps to open a VNC connection:

1. Open the VNC Web browser by typing the following commands:

```
# cd /root
# vncserver
```

2. When you are asked for a password, type a new password and write it down. A message displays a session number, as in the following example:

```
hostIP:1
```

where 1 is the session number.

3. Open a local browser window and enter *<host IP:580 session number>*, as shown in the following example:

```
http://9.27.95.21:5809
```

where the hostIP is 9.27.95.21, and the session number is 9.

Steps for installing DS4000 Storage Manager software using the installation wizard: Complete the following steps to install the DS4000 Storage Manager software using the installation wizard. Adjust the steps as necessary for your specific installation.

1. Download the Storage Manager 9.14 SMIA package file from the DS4000 Storage Manager CD, or from the DS4000 support Web site, to a directory on your system by typing the following command:

```
sh SMIA-LINUXPPC-09.14.A5.05.bin
```

2. Open the file on your system. The Storage Manager Installation wizard's Introduction window opens.

3. Follow the instructions in each window of the wizard. When you select the installation type, you can choose one of the following options:
   - Typical (Full) Installation — Installs all Storage Manager software packages
   - Management Station — Installs SMruntime and SMclient
   - Host — Installs SMruntime, SMagent and SMutil
   - Custom — Allows you to select which packages you want to install

Result: The DS4000 Storage Manager software is installed on your system.

Installing Storage Manager packages manually

Use the following procedures, in the order listed, to install the DS4000 Storage Manager software on a POWER-based RHEL 3 or SLES 9 storage management station.

Install the Storage Manager software in the following order.
1. "Steps for installing SMruntime software" on page 148
2. "Steps for installing SMclient software" on page 148
3. "Steps for installing SMutil" on page 149
4. "Steps for installing SMagent (optional)" on page 149

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 RHEL 3 or SLES 9 host on which you are installing SMruntime meets the minimum hardware and software requirements described in "Hardware, firmware and system requirements" on page 133.
- The Graphical Interface X Configuration is installed and configured. If this was not completed at Linux installation, configure it now using one of the following tools:

  RHEL 3   Use `redhat-config-xfree86`.

  SLES 9   Use `xfree86`.

- The DS4000 Storage Manager files have been downloaded to a directory on the host.

  Note: If you do not have the files, you can find them at the following Web site:

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

  Click the link for your storage subsystem. When the storage subsystem page opens, click the **Download** tab, and follow the links to the Storage Manager files.

Steps for installing SMruntime software: Modify the following commands as needed for your installation. No restart is required during the installation process.

1. Install SMruntime by typing the following command:

```
# rpm -ihv SMruntime-LINUX-09.14.A5.00-1.ppc64.rpm
```

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

```
# rpm -qa | grep SMruntime
```

Steps for installing SMclient software: Modify the following commands as needed for your installation. No restart is required during the client installation process.

1. Install SMclient by typing the following command:

```
# rpm -ihv SMclient-LINUX-09.14.G5.05-1.noarch.rpm
```

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

```
# rpm -qa |grep SMclient
```

   Result: The verification process shows the package name and version.

After completing the installation of the client software, you can launch the SMclient on the host by typing the following command:

```
# /opt/IBM_DS4000/SMclient
```

Result: The Storage Manager Enterprise Management window opens.

Note: If the Enterprise Management window does not open, verify that the Storage Manager software packages listed in the installation procedures completed their installation without error, and that the X Graphical Interface is configured and running on the server.

(The SMclient is a Java-based program that must run in a graphical interface or remote graphical window. See "Opening a remote Xwindows or VNC connection" on page 146.)

If the verification process returns an error, contact your support representative.

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

1. Install SMutil by typing the following command:

```
# rpm -ihv SMutil-LINUX-09.14.A5.00-1.ppc64.rpm
```

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

```
# rpm -qa |grep SMutil
```

   Result: The verification process shows the package name and version.

Steps for installing SMagent (optional): SMagent is required for in-band management only. It is not required for out-of-band management.

Note: SMagent is not supported with DS4000 controller firmware version 5.42.xx.xx.

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

1. Install SMagent by typing the following command:

```
# rpm -ihv SMagent-LINUX-09.14.A5.00-1.ppc64.rpm
```

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

```
# rpm -qa |grep SMagent
```

   Result: The verification process shows the package name and version.

Note: If no UTM LUNs are assigned on the DS4000, SMagent stops with a warning message that there are no UTM LUNs available. This is normal. After a UTM LUN is assigned, start the SMagent program by typing `SMagent start`.

Configuring storage subsystems

Before you can configure storage subsystems, the subsystem must be physically configured for direct management using Ethernet connections from the management station to each controller. Also, SMclient must be installed on the management station before you can configure the subsystem.

Notes:

1. See "Setting up IP addresses for DS4000 storage controllers" on page 20 for information about assigning initial IP addresses to the controllers.
2. If you are using a RHEL 3 or SLES 9 host that is a headless server, you can configure DS4000 Storage Manager by remote vnc connection or by using SMclient on a different management station.

This section contains the following topics:

- "Adding storage subsystems to SMclient"
- "Storage Partitioning: Defining a host group" on page 155

Adding storage subsystems to SMclient

After the disk subsystem is configured on the network, perform the following steps to add storage subsystems to DS4000 SMclient.

Notes:

1. If the Enterprise Management window does not open, verify that the Storage Manager software packages listed in the installation procedures completed their installation without error, and that the X Graphical Interface is configured and running on the server.

   (The SMclient is a Java-based program that must run in a graphical interface or remote graphical window. See "Opening a remote Xwindows or VNC connection" on page 146.)

   If the verification process returns an error, contact your support representative.

2. The Storage Manager client software version 9.12, and DS4000 controller firmware 06.12.xx.xx, are not supported on SLES 9. You need a different type of management station to manage the storage subsystems of you are using Storage Manager 9.12 or DS4000 controller firmware 06.12.xx.xx with a SLES 9 host.

3. The Storage Manager client software version 9.12 is supported on RHEL 3, and Storage Manager client software version 9.14 is supported on both SLES 9 and RHEL 3, but you can use a different type of management station to manage the storage subsystems if you prefer.

1. Start SMclient on the management station by typing the following command:

```
# /opt/IBM_DS4000/SMclient
```

The Enterprise Management window opens.

2. Complete the following steps to add the storage subsystems to the SMclient:

   a. In the Enterprise Management window, click **Edit —> Add Storage Subsystem**.

   b. In the Add Storage Subsystem window, type the IP addresses of each controller in the storage subsystem, and click **Add**. The name of the storage subsystem displays in the Enterprise Management window.

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

3. Double-click the name of the storage subsystem to open its Subsystem Management window.

Upgrading controller firmware and NVSRAM

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

Overview

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

- "Using concurrent download" on page 151
- "Upgrading controller firmware and NVSRAM"
- "Using traditional or staged download" on page 151

Using concurrent download: Depending on your system's current firmware version, 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.

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 DS4000 controller firmware version 5.4x.xx.xx to version 6.1x.xx.xx. You can only use staged download if the DS4000 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 subsystems.

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 subsystem 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 subsystem that you plan to update.

Preparing for firmware or NVSRAM download

To prepare for download, review and complete all of the steps in this section before you start the download procedures.

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

To check the controller 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 view the current versions of controller firmware and NVSRAM.
3. Click **Close** to close the Storage Subsystem Profile window.
4. To find whether the versions of controller firmware and NVSRAM that are currently on your storage subsystem are the latest versions, see the Storage Manager readme file for Linux on POWER at the following Web site:

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

   Click the link for your storage subsystem. When the storage subsystem page opens, click the **Download** tab, and follow the links to the readmes.
5. If you find that you do not have the latest versions of the firmware and NVSRAM, download the correct versions from this Web site to a directory on your system, then follow the preparation and upgrade procedures that are described in this section.

Save the storage subsystem configuration: Saving the storage subsystem configuration to a script file is necessary in order to restore the configuration in the event of a catastrophic failure.

Attention: Do not use this option if the storage subsystem is undergoing an operation that changes any configuration settings. (If any logical drive listed in the Logical View of the Subsystem Management window has a clock icon, then a configuration change is occurring, and you must wait to save the configuration.)

You can also save the storage subsystem profile by selecting **Storage Subsystem —> View Profile** in the Subsystem Management window, and clicking **Save As**. Then, in the Save Profile window, select the sections of the profile that you want to save, specify a name and location for the text file, and click **Save**.

Save the storage subsystem configuration by completing the following steps:
1. In the Subsystem Management window, click **Storage Subsystem —> Configuration —> Save**. The Save Configuration window displays.
2. In the Save Configuration window, select the elements of the configuration that you want to save, and click **Yes**. A save dialog displays.
3. In the save dialog, specify a name for the script file and select its location.

   Notes:
   - The .cfg extension is used by default for the script file if you do not specify a file extension.
   - Do not store your DS4000 Storage Subsystem configuration script file in the same location as the logical drives that are defined on your DS4000 Storage Subsystem. If you do, and the logical drives fail during the subsequent migration, you will lose data.
4. Click **Save**. The Configuration Saved window displays.
5. Review the information in the Configuration Saved window, then click **OK** to return to the Subsystem Management window.

Determine whether you can use concurrent download: You can use concurrent download only if you are upgrading from one version of controller firmware to a later release of the same version. For example, if you are upgrading from 06.10.xx.xx to 6.12.xx.xx, you can use concurrent download. However, if you are upgrading from 5.4x.xx.xx to 06.1x.xx.xx, you *cannot* use concurrent download, and you need to stop I/O before upgrading controller firmware.

Note: You can only use concurrent download if an RDAC driver is installed on the host.

Table 29. Firmware versions that support concurrent download (Linux on POWER)

DS4000 storage subsystem:	Controller firmware version:
DS4300 base and Turbo:	6.10.xx.xx. 6.12.xx.xx
DS4300 single-controller (SCU):	05.34.07.xx
DS4400/DS4500:	6.10.xx.xx, 6.12.xx.xx
DS4800	06.14.xx.xx
DS4100	06.10.xx.xx. 06.12.xx.xx
All other DS4000 storage subsystems:	05.30.xx.xx, 6.10.xx.xx, 6.12.xx.xx

Steps for upgrading the controller firmware

Complete the following steps to install a new firmware version on the controllers:

1. Download the correct version of the firmware from the Web site that is listed in "Check current versions of controller firmware and NVSRAM" on page 151. Place the file in a designated directory on the host system.

2. Check your system's current firmware versions to determine whether you need to stop I/O, using the information that is listed in Table 29. In addition, verify that an RDAC driver is installed on the host.

   Attention: If you do not have the correct firmware versions for concurrent download, or if an RDAC driver is not installed, you *must* stop I/O to the controllers before you upgrade the firmware.

3. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.

4. In the Subsystem Management window, click **Advanced —> Maintenance —> Download —> Controller Firmware**. The Download Firmware window displays.

5. Click **Browse** to select the firmware file that you previously downloaded to a directory on the host.

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**. The Confirm Download window displays.

8. Click **Yes** to start the download.

9. Close the Subsystem Management window and then reopen it to complete the firmware update.

10. Verify that the firmware was successfully installed.

Steps for upgrading the NVSRAM

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

1. Download the correct NVSRAM version from the Web site that is listed in "Check current versions of controller firmware and NVSRAM" on page 151. Place the file in a designated directory on the host system.
2. Stop I/O.
3. Ensure that the array is in an Optimal state. If it is not, resolve the issue before continuing to the next step.
4. In the Subsystem Management window, click **Advanced —> Maintenance —> Download —> Controller NVSRAM**. The Download NVSRAM window displays.
5. Click **Browse** to select the NVSRAM file that you previously downloaded to a directory on the host.
6. Click **OK**. The Confirm Download window displays.
7. Click **Yes** to start the download.
8. Verify that the NVSRAM was successfully installed.
9. Start I/O.

Creating direct-attached and SAN-attached configurations

DS4000 Storage Manager supports IBM DS4000 Storage Subsystems 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

The following types of direct-attached configurations are supported:

- One to four POWER-based Linux host servers, in single-path configurations, can be connected to each DS4000 storage subsystem.
- One or two POWER-based Linux servers, in dual-path or RDAC configurations, can be connected to each DS4000 storage subsystem. Each server must have exactly one HBA connected to each controller port.
- Single and dual HBA configurations are supported.

Note: Two-server, dual-path DS4400 or DS4500 configurations require four host-side minihubs, each with exactly one Fibre Channel connection from each HBA to a minihub.

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

1. Ensure that no external hubs are being used.
2. Connect the HBAs to each controller or minihub port of the DS4000 storage subsystem.
3. Configure and verify the configuration.

Creating a SAN-attached configuration

Before you begin: 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 in a host system must be isolated from each other if they are connected to the same switch that is connected to the same DS4000 controller port.
- An RDAC configuration requires two HBAs per host system, each in a separate switch zone that is connected to one of the two controllers.
- Each Fibre Channel HBA on each host, and controller port on each DS4000 subsystem, must be in its own fabric zone, if they are connecting through a single 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.
4. Configure, then verify the configuration by logging into the switch to view the system.

Note: For more information about zoning and enabling zones, see Appendix F, "Connecting HBAs in a FC switch environment," on page 223.

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.

If you have purchased a Storage Partitioning premium feature, ensure that the premium feature is enabled. For more information about storage partitioning, see "Storage Partitioning premium feature" on page 24.

Steps for defining a host group

Complete the following steps to define a host group.

Before you begin: Note the following information.

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

  Exception: If you are running a DS4000 configuration without partitioning enabled, you can use the default host group.

- Auto volume transfer/auto-disk transfer (AVT/ADT) is managed by the Linux RDAC driver. To avoid contention, you must disable AVT/ADT by selecting the `LNXCL` host type instead of the `Linux` host type when you define your host groups, as described in the procedure. (The `Linux` host type automatically enables AVT/ADT; the `LNXCL` host type does not.)

1. Click the **Mappings View** tab on the Subsystem Management window.
2. In the Topology section of the Mappings window, highlight the name of the storage subsystem, and click **Mappings —> Define —> Host Group**.

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

3. Type the name of the new host group (for example, `RHEL 3`). 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, `redhatserver`). 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, as follows, according to your configuration:
   - **Single-path:** Select **Linux**.
   - **Dual-path:** Select **LNXCL**.

   Click **Add**.

   Important: Failure to change the host type from the default might cause undesired results.

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

   Note: You must configure two HBAs per partition if you are using RDAC.

   Important: The host port identifier is the worldwide port name (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.

   Note: For more information about zoning and enabling zones, see Appendix F, "Connecting HBAs in a FC switch environment," on page 223.

   lpfc module
   You might not have initialized the lpfc module (verify with the **lsmod** command). Type **modprobe lpfc** or **modprobe lpfcdfc** for the subsystem that does not display in the menu.

Steps for verifying that the host ports match the host

1. Type the following command to view the HBAs:

```
# ls /proc/scsi/lpfc
```

   Result: A list containing all of the HBAs discovered on the system is shown.

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/lpfc#
```

   where *lpfc#* is the assigned HBA number.

   Result: Information about the HBA displays, as shown in the following example. Note that the Portname value (**bold** text) is equivalent to the WWPN:

```
# cat /proc/scsi/lpfc/3
Emulex LightPulse FC SCSI 7.0.3
Emulex LightPulse LP9802 2 Gigabit PCI Fibre Channel
 Adapter on PCI
bus 161 device 08 irq 137
SerialNum: 0000c930ae1e
Firmware Version: 1.00 (H2J1.00X5)
Hdw: 2003806d
VendorId: 0xf98010df
Portname: 10:00:00:00:c9:30:ae:1e   Nodename: 20:00:0
0:00:c9:30:ae:1e

Link Up - Ready:
   PortID 0x11400
   Fabric
   Current speed 2G
```

4. Verify that the WWPN matches the Host Port Identifier number assigned to the host port in the host partition table of the DS4000 SMclient.

5. Repeat this procedure to verify the second host port, if configuring for RDAC.

Mapping LUNS to a storage partition

This section contains the following procedures:

- "Mapping LUNs to a new partition"
- "Adding LUNs to an existing partition" on page 158

Important:

1. It is recommended that you always map the LUNs using consecutive numbers, starting with LUN 0. For example, map LUNs to numbers 0; 1; 2; 3; 4; 5; and so on, without skipping any numbers.

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

3. On each partition, you must map a LUN 0. (It is recommended that LUN 0 is not a UTM LUN.)

4. You must map at least one non-UTM LUN on each partition.

Mapping LUNs to a new partition

When mapping LUNs for the first time to a newly created partition, complete the following steps to map LUNs to the partition.

Before you begin: Read "DS4000 Storage Manager performance planning" on page 21 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 section, right-click the host on which you want to map LUNs, and select **Define Storage Partitioning**. The Define Storage Partitioning window opens.

3. In the Define Storage Partitioning window, select **Host**, then click **Next**.

4. Select the logical drive by name, on the right side of the window.

5. Accept the default LUN ID, or change it, then click **Add**.

6. Repeat step 5 for each LUN that you want to map to the partition.

Adding LUNs to an existing partition

When adding LUNs to an existing partition, complete the following steps to map the new LUNs to the partition.

Before you begin: Read "DS4000 Storage Manager performance planning" on page 21 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 section, 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 that you want to add to the partition.

Configuring the host bus adapters

Complete the procedure that is appropriate for your operating system:

- "Configuring SLES 9 host bus adapters"
- "Configuring RHEL 3 host bus adapters" on page 160

Configuring SLES 9 host bus adapters

Each Emulex HBA that is used with POWER-based SLES 9 hosts needs two separate HBA and utility driver packages. The standard packages are loaded onto the system with the SLES 9 installation.

Use the **lsmod** command to find out whether the HBA device driver (lpfcdd module) is on the system.

Configuring HBAs: Complete the following steps to set up the HBAs:

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:

```
hostname:~ # 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
```

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.

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 HBA properties" on page 169.)

5. Type the following command to display all assigned LUNs:

```
# SMdevices
```

Result: A list of information about all assigned LUNs is displayed. This information includes the assigned host device name (/dev/sd..) and the following DS4000 storage subsystem information about the owning storage subsystem, as shown by the SMclient:

- Name
- Assigned Logical drive name
- Assigned LUN ID
- Preferred controller owner, and whether that controller is currently controlling the logical drive

The output of **SMdevices** command looks similar to the following example:

```
# SMdevices
IBM FAStT Storage Manager Devices, Version 09.12.A5.00
Built Fri Jan 14 16:42:15 CST 2005
(C) Copyright International Business Machines Corporation,
2004 Licensed Material - Program Property of IBM.  All rights reserved.

  /dev/sdh (/dev/sg10) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-5-0A, LUN 4, Logical Drive ID
<600a0b80000f0fc300000044412e2dbf>, Preferred Path (Controller-A): In Use]
  /dev/sdd (/dev/sg6) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-5-1A, LUN 0, Logical Drive ID
<600a0b80000f13ec00000016412e2e86>, Preferred Path (Controller-B): In Use]
  /dev/sde (/dev/sg7) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-0-0A, LUN 1, Logical Drive ID
<600a0b80000f0fc30000003c412e2d59>, Preferred Path (Controller-A): In Use]
  /dev/sdf (/dev/sg8) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-1-0A, LUN 2, Logical Drive ID
<600a0b80000f0fc30000003e412e2d79>, Preferred Path (Controller-A): In Use]
  /dev/sdg (/dev/sg9) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-3-0A, LUN 3, Logical Drive ID
<600a0b80000f13ec00000012412e2e4c>, Preferred Path (Controller-A): In Use]
```

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.

Configuring RHEL 3 host bus adapters

Complete the following steps to set up HBAs on a RHEL 3 host.

1. Type the following command to verify that the lpfc and lpfcdfc modules are installed and working:

```
# lsmod
```

   Result: A list of installed modules displays.

2. Verify that entries for the lpfc and lpfcdfc modules are included in the **lsmod** list.
3. Create the SAN switch connections.
4. Use one of the following methods to probe for the HBA:

   • **Method 1: rmmod, modprobe**

   Type the following commands:

With Emulex Kit installed	Without Emulex Kit installed
# rmmod lpfcdfc # rmmod lpfc # modprobe lpfcdfc	# rmmod lpfc # modprobe lpfc

   Result: The driver searches for all HBAs, and their connections, that are on the system.

   • **Method 2: lun_scan**

   Type the following command:

```
# cd /usr/sbin/lpfc
# lun_scan hba#
```

   to rescan a particular HBA if the Applications package is installed, or

```
# cd /usr/sbin/lpfc
# lun_scan all
```

to rescan all the HBAs that are on the system.

Result: The driver searches for the HBA or HBAs, and the HBA connections.

5. 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 HBA properties" on page 169.)

6. Type the following command to display all assigned LUNs:

```
# SMdevices
```

Result: A list of information about all assigned LUNs is displayed. This information includes the assigned host device name (/dev/sd..) and the following DS4000 storage subsystem information about the owning storage subsystem, as shown by the SMclient:

• Name

• Assigned Logical drive name

• Assigned LUN ID

• Preferred controller owner, and whether that controller is currently controlling the logical drive

The output of **SMdevices** command looks similar to the following example:

```
# SMdevices
IBM FAStT Storage Manager Devices, Version 09.12.A5.00
Built Fri Jan 14 16:42:15 CST 2005
(C) Copyright International Business Machines Corporation,
2004 Licensed Material - Program Property of IBM.  All rights reserved.

  /dev/sdh (/dev/sg10) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-5-0A, LUN 4, Logical Drive ID
<600a0b80000f0fc300000044412e2dbf>, Preferred Path (Controller-A): In Use]
  /dev/sdd (/dev/sg6) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-5-1A, LUN 0, Logical Drive ID
<600a0b80000f13ec00000016412e2e86>, Preferred Path (Controller-B): In Use]
  /dev/sde (/dev/sg7) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-0-0A, LUN 1, Logical Drive ID
<600a0b80000f0fc30000003c412e2d59>, Preferred Path (Controller-A): In Use]
  /dev/sdf (/dev/sg8) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-1-0A, LUN 2, Logical Drive ID
<600a0b80000f0fc30000003e412e2d79>, Preferred Path (Controller-A): In Use]
  /dev/sdg (/dev/sg9) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-3-0A, LUN 3, Logical Drive ID
<600a0b80000f13ec00000012412e2e4c>, Preferred Path (Controller-A): In Use]
```

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.

Installing the RDAC Multipath Proxy driver

This section describes how to install the RDAC driver for a dual-path POWER-based Linux configuration.

Important: Before you install RDAC, make sure that the partitions and LUNs are configured and assigned.

Complete the following steps to install RDAC:

1. Download the RDAC driver package from the IBM DS4000 support Web site.

   RDAC package for RHEL 3:
   rdac-LINUX-09.00.A5.09-source.tar.gz

   RDAC package for SLES 9:
   rdac_LINUX_09.01.B5.02_source_tar_gz

2. Create a directory on the host and download the RDAC driver package to that directory.

3. Uncompress the file by typing the following command:

```
# tar -zxvf rdac-LINUX-package_version-source.tar.gz
```

   where *package_version* is the RHEL 3 or SLES 9 package version number that is specified in step 1.

   Result: A directory called **linuxrdac** is created.

4. Open the Readme.txt that is included in the linuxrdac directory.

5. In the Readme.txt, find the instructions for building and installing the driver and complete the steps, which include rebooting the server.

   Note: Make sure you reboot the server before you proceed to step 6.

6. Type the following command to list the installed modules:

```
# lsmod
```

7. Verify that module entries are included in the **lsmod** list, as follows:

   Module entries for RHEL 3
   - scsi_mod
   - sd_mod
   - sg
   - lpfc
   - mpp_Vhba
   - mpp_Upper
   - lpfcdfc (if ioctl module is installed)

   Module entries for SLES 9
   - mppVhba
   - mppUpper
   - lpfcdd

   Note: If you do not see the mpp_Vhba module, the likely cause is that the server was rebooted before the LUNs were assigned, so the mpp_Vhba module was not installed. If this is the case, assign the LUNs now, reboot the server, and then repeat step 7.

8. Type the following command to verify the driver version:

```
# mppUtil -V
```

   Result: The Linux multipath driver version displays.

9. Type the following command to verify that devices are configured with the RDAC driver

```
# ls -1R /proc/mpp
```

Result: An output similar to the following example displays:

```
# ls -1R /proc/mpp
/proc/mpp:
total 0
dr-xr-xr-x    4 root      root                0 Oct 24 02:56 DS4100-sys1
crwxrwxrwx    1 root      root        254,    0 Oct 24 02:56 mppVBusNode

/proc/mpp/ DS4100-sys1:
total 0
dr-xr-xr-x    3 root      root                0 Oct 24 02:56 controllerA
dr-xr-xr-x    3 root      root                0 Oct 24 02:56 controllerB
-rw-r--r--    1 root      root                0 Oct 24 02:56 virtualLun0
-rw-r--r--    1 root      root                0 Oct 24 02:56 virtualLun1
-rw-r--r--    1 root      root                0 Oct 24 02:56 virtualLun2
-rw-r--r--    1 root      root                0 Oct 24 02:56 virtualLun3
-rw-r--r--    1 root      root                0 Oct 24 02:56 virtualLun4
-rw-r--r--    1 root      root                0 Oct 24 02:56 virtualLun5

/proc/mpp/ DS4100-sys1/controllerA:
total 0
dr-xr-xr-x    2 root      root                0 Oct 24 02:56 lpfc_h6c0t2

/proc/mpp/ DS4100-sys1/controllerA/lpfc_h6c0t2:
total 0
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN0
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN1
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN2
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN3
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN4
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN5

/proc/mpp/ DS4100-sys1/controllerB:
total 0
dr-xr-xr-x    2 root      root                0 Oct 24 02:56 lpfc_h5c0t0

/proc/mpp/ DS4100-sys1/controllerB/lpfc_h5c0t0:
total 0
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN0
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN1
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN2
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN3
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN4
-rw-r--r--    1 root      root                0 Oct 24 02:56 LUN5
```

Note: After you install the RDAC driver, the following commands and man pages are available:
- mppUtil
- mppBusRescan
- mppUpdate

Performing initial device identification

Complete the procedure that is appropriate for your operating system:
- "Performing initial device identification on SLES 9 hosts" on page 164
- "Performing initial device identification on RHEL 3 hosts" on page 166

Performing initial device identification on SLES 9 hosts

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:

```
hostname:~ # 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
```

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.

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:

```
hostname:/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
hostname:/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:

```
hostname:/sys/class/scsi_host/host1 # ls
.   cmd_per_lun  host_busy  scan         unchecked_isa_dma
..  device       proc_name  sg_tablesize unique_id

hostname:/sys/class/scsi_host/host1 # cd device
hostname:/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 **lsscsi** 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 **lsscsi** 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. Use the browser to log into the switch.

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**.

A report showing port status, zones and connections is displayed.

Performing initial device identification on RHEL 3 hosts

Type the following command to view the LUNs that the host can detect:

```
# cat /proc/scsi/scsi
```

The resulting output depends on whether you have RDAC installed.

Device identification without RDAC installed

If you do not have the RDAC driver installed, you see output similar to the following example:

```
# cat /proc/scsi/scsi

Attached devices:
Host: scsi0 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: DROM00203        Rev: 7R41
  Type:   CD-ROM                           ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 08 Lun: 00
  Vendor: IBM      Model: ST336607LC       Rev: C50F
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 1742-900         Rev: 0520
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: 1742-900         Rev: 0520
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 02
  Vendor: IBM      Model: 1742-900         Rev: 0520
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 03
  Vendor: IBM      Model: 1742-900         Rev: 0520
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 04
  Vendor: IBM      Model: 1742-900         Rev: 0520
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 05
  Vendor: IBM      Model: 1742-900         Rev: 0520
  Type:   Direct-Access                    ANSI SCSI revision: 03
```

Note: If two paths are configured, you will see each LUN two times — once on each configured HBA path.

If the output shows that no LUNs are assigned to the host, reboot the host or type the following sequence of commands to probe the scsi bus for the added devices:

```
# rmmod lpfcdfc
# rmmod lpfc
# modprobe lpfcdfc
```

Then, type the following command again:

```
# cat /proc/scsi/scsi
```

Device identification with RDAC installed

If you have the RDAC driver installed, you see output similar to the following sample output. In the output, there are five LUNs from a DS4500 storage subsystem (Model: 1742-900), assigned to two HBAs:

```
# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: DROM00203       Rev: 7R41
  Type:   CD-ROM                          ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 08 Lun: 00
  Vendor: IBM      Model: ST336607LC      Rev: C50F
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 1742-900        Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: 1742-900        Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 02
  Vendor: IBM      Model: 1742-900        Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 03
  Vendor: IBM      Model: 1742-900        Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 04
  Vendor: IBM      Model: 1742-900        Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 1742-900        Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: 1742-900        Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 02
  Vendor: IBM      Model: 1742-900        Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 03
  Vendor: IBM      Model: 1742-900        Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 04
  Vendor: IBM      Model: 1742-900        Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi11 Channel: 00 Id: 04 Lun: 00
  Vendor: IBM      Model: VirtualDisk     Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi11 Channel: 00 Id: 04 Lun: 01
  Vendor: IBM      Model: VirtualDisk     Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi11 Channel: 00 Id: 04 Lun: 02
  Vendor: IBM      Model: VirtualDisk     Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi11 Channel: 00 Id: 04 Lun: 03
  Vendor: IBM      Model: VirtualDisk     Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi11 Channel: 00 Id: 04 Lun: 04
  Vendor: IBM      Model: VirtualDisk     Rev: 0520
  Type:   Direct-Access                   ANSI SCSI revision: 03
```

Notes:

1. Each LUN is displayed on both of its configured paths.

2. The **VirtualDisk** is the RDAC's logical disk. There is one VirtualDisk for each LUN, regardless of number of paths available.

3. If you do not see any VirtualDisks in the output, the most likely cause is that the mpp_Vhba module did not get installed. (This happens if the RDAC driver was installed before any LUNs were configured.) Use the **lsmod** command to verify whether the mpp_Vhba module is installed. If it is missing from the list of modules, make sure that the LUNs are configured, then reboot the host.

Performing common maintenance tasks

See the information in this section that is appropriate for your operating system:
- "Performing SLES 9 maintenance tasks"
- "Performing RHEL 3 maintenance tasks" on page 172

Performing SLES 9 maintenance tasks

This section describes some common system administration tasks.
- "Viewing HBA properties"
- "Checking LUN size" on page 171

Viewing HBA properties

You can use one of the following methods to view HBA properties, such as WWPNs and firmware levels.

Viewing HBA properties, 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:

```
hostname:/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
hostname:/proc/scsi/lpfc #
```

Viewing HBA properties, 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
```

Checking LUN size

Complete the following steps to check the size of a LUN:

1. From the device directory, type the **cat size** command, as shown in the following example:

```
hostname:/sys/class/scsi_host/host11/device/11:0:0:0/block # cat size
```

Result: A number displays, as in the following example:

```
8388608
```

2. Multiply this number by 512 (bytes) to calculate the size of the LUN, as shown in the following example:

```
8388608 * 512 = 4294967296 (~ 4GB)
```

Result: The result of the calculation is the size of the LUN. In the example, the LUN size is approximately 4 GB.

Using dynamic capacity expansion and dynamic volume expansion

Dynamic volume expansion (DVE) increases the size of a logical drive. In order to perform a DVE, there must be free capacity available on the array. If there is not, you can first perform a dynamic capacity expansion (DCE) to increases the array's capacity by adding drives.

Performing a dynamic capacity expansion operation:

Before you begin: You can find more information about this procedure in the Storage Manager online help.

Complete the following steps to increase the capacity on the array by performing a DCE:

1. In the Logical/Physical view of the Subsystem Management window, right-click on an array and select **Add Free Capacity (Drives)**.

2. In the Add Free Capacity (Drives) window, select one or two available drives and click **Add**.

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

Performing a dynamic volume expansion operation:

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, you can add one or more to the array by performing a dynamic capacity expansion (DCE) operation before you perform the DVE operation.

You can find more information about this procedure in the Storage Manager online help.

Restrictions:

1. You cannot resize the logical drive while the logical drive group is activated in classic or enhanced concurrent mode.
2. You cannot resize the root logical drive group.
3. Dynamic Volume Expansion is not supported on 2.4 kernels.

Complete the following steps to increase the size of a logical drive by performing a DVE:

1. Stop I/O to the logical drive whose size you want to increase.
2. Unmount the logical drive, if it is mounted.
3. From the Logical/Physical window of the Subsystem Management window, right-click the logical drive and select **Increase Capacity**. The Increase Logical Drive Capacity — Additional Instructions window opens.
4. Read the additional instructions and click **OK**. The Increase Logical Drive Capacity window opens.
5. Type the amount that you want to increase the logical drive, and click **OK**.

   You see a clock icon on every logical drive within the array. You must wait for the process to complete before you can begin any host intervention.

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

6. On the host, rescan the logical drive by typing the following command from the device directory:

```
# echo 1 rescan
```

7. Check the size of the logical drive using the steps that are described in "Checking LUN size" on page 171.
8. Remount the logical drive.

Performing RHEL 3 maintenance tasks

This section describes some common system administration tasks.

Using the SMdevices utility to find information about LUNs

SMutil features an excellent utility, **SMdevices**, which you can use to map the host-assigned device name for each LUN back to its corresponding DS4000 storage subsystem device.

In the **SMdevices** output, you can view the following DS4000 storage subsystem information, as it is shown on SMclient.

Note: The examples in the list refer to the sample **SMdevices** output.

- Host assigned name (`/dev/sdh`)
- DS4000 storage subsystem name (`DS4500_Storage_Server-A`)
- Logical drive name `Raid-5-0A`
- LUN ID (`LUN 4`)
- Preferred controller owner, and whether that controller is currently controlling the logical drive

The following example shows a sample **SMdevices** output for the subsystem **DS4500_Storage_Server-A**:

```
# SMdevices
IBM FAStT Storage Manager Devices, Version 09.12.A5.00
Built Fri Jan 14 16:42:15 CST 2005
(C) Copyright International Business Machines Corporation,
2004 Licensed Material - Program Property of IBM.  All rights reserved.

  /dev/sdh (/dev/sg10) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-5-0A, LUN 4, Logical Drive ID
<600a0b80000f0fc300000044412e2dbf>, Preferred Path (Controller-A): In Use]
  /dev/sdd (/dev/sg6) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-5-1A, LUN 0, Logical Drive ID
<600a0b80000f13ec00000016412e2e86>, Preferred Path (Controller-B): In Use]
  /dev/sde (/dev/sg7) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-0-0A, LUN 1, Logical Drive ID
<600a0b80000f0fc30000003c412e2d59>, Preferred Path (Controller-A): In Use]
  /dev/sdf (/dev/sg8) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-1-0A, LUN 2, Logical Drive ID
<600a0b80000f0fc30000003e412e2d79>, Preferred Path (Controller-A): In Use]
  /dev/sdg (/dev/sg9) [Storage Subsystem DS4500_Storage_Server-A,
Logical Drive Raid-3-0A, LUN 3, Logical Drive ID
<600a0b80000f13ec00000012412e2e4c>, Preferred Path (Controller-A): In Use]
```

Viewing queue depth

To view the queue depth of each LUN, type the following command:

```
# cat /proc/scsi/sg/device_hdr devices
```

Note: If RDAC is installed, the queue depth number that is displayed in the output is not the same as the host setting for queue depth, because the RDAC driver takes two queues for each LUN. For example, in the following sample output, the queue depth of LUNs mapped to host 11 is set to 10; because RDAC is installed, the output shows a queue depth of 8.

```
#cat /proc/scsi/sg/device_hdr devices
host    chan    id    lun    type    opens    qdepth    busy    online
0       0       1     0      5       0        2         0       1
0       0       8     0      0       3        8         0       1
11      0       4     0      0       0        8         0       1
11      0       4     1      0       0        8         0       1
11      0       4     2      0       0        8         0       1
11      0       4     3      0       0        8         0       1
11      0       4     4      0       0        8         0       1
```

Replacing IBM host bus adapters using PCI Hotplug tools

This section provides requirements and procedures for replacing IBM host bus adapters in pSeries servers on your SLES 9 SP1 host system using PCI Hotplug tools.

Requirements

Operating system

SUSE Linux Enterprise Server 9 Service Pack 1 (SLES 9 SP1)

Kernel: 2.6.5-7.139-pseries

PCI Hotplug tools

Ensure that the following tools are installed in the /usr/sbin directory:

- **lsslot**
- **drslot_chrp_pci**

If these tools are not installed, complete the following steps to install them:

1. Ensure that `rdist-6.1.5-792.1` and `compat-2004.7.1-1.2` are installed from the SLES 9 media.
2. Download the PCI Hotplug Tools rpm files from the following Web site:

   techsupport.services.ibm.com/server/lopdiags/suselinux/hmcserver
3. On the Web page, select **SLES 9 Tools**. Download and install the following rpm files:
   - librtas-1.1-17.ppc64.rpm
   - rpa-pci-hotplug-1.0-10.ppc.rpm
4. Type the following command to install each rpm file:

```
# rpm -Uvh <filename>.rpm
```

where *<filename>* is the name of the rpm file.

Host bus adapters

The following host bus adapters are supported with this procedure:

- FC 6228
- FC 6239
- FC 5716

PCI core

The PCI core must be loaded on the system. Type the following command to verify:

```
# ls -l /sys/bus/pci/slots
```

If the /sys/bus/pci/slots directory does not exist, then the PCI core is not loaded.

rpaphp driver

The rpaphp driver must be loaded on the system. Type the following command to verify:

```
ls -l /sys/bus/pci/slots/*
```

Listing information about the I/O slots

Before you replace an HBA using PCI Hotplug, you can use the **lsslot** tool to list information about the I/O slots. This section describes how to use **lsslot**, and provides examples.

Syntax for the lsslot command: Use the **lsslot** tool as follows:

- Syntax: **lsslot [-c slot | -c pci [-a | -o]] [-s drc-name] [-F delimiter]**
- Options:

  No options   Displays all DR slots

  -c slot   Displays all DR slots

  -c pci   Displays all PCI Hotplug slots

  -c pci -a   Displays all available (empty) PCI Hotplug slots

  -c pci -o   Displays all occupied PCI Hotplug slots

  -F   Uses delimiter to delimit columns

Listing PCI Hotplug slots using the lsslot command: This section shows the command lines you can use to list PCI Hotplug slots.

Note: In the Device(s) columns of the command-line outputs, the PCI devices in the slots are listed as follows: *xxxx:yy:zz.t*. (For example: 0001:58:01.1)

List all PCI Hotplug slots: Type the following command to list all PCI Hotplug slots:

```
# lsslot -c pci
```

The resulting output looks similar to the following:

```
# Slot                  Description                       Device(s)
U7879.001.DQD014E-P1-C1  PCI-X capable, 64 bit, 133MHz slot  Empty
U7879.001.DQD014E-P1-C2  PCI-X capable, 64 bit, 133MHz slot  0002:58:01.0
U7879.001.DQD014E-P1-C3  PCI-X capable, 64 bit, 133MHz slot  0001:40:01.0
U7879.001.DQD014E-P1-C4  PCI-X capable, 64 bit, 133MHz slot  Empty
U7879.001.DQD014E-P1-C5  PCI-X capable, 64 bit, 133MHz slot  Empty
U7879.001.DQD014E-P1-C6  PCI-X capable, 64 bit, 133MHz slot  0001:58:01.0
0001:58:01.1
```

List all occupied PCI Hotplug slots: Type the following command to list all occupied PCI Hotplug slots:

```
# lsslot -c pci -o
```

The resulting output looks similar to the following:

```
# Slot                  Description                       Device(s)
U7879.001.DQD014E-P1-C2  PCI-X capable, 64 bit, 133MHz slot  0002:58:01.0
U7879.001.DQD014E-P1-C3  PCI-X capable, 64 bit, 133MHz slot  0001:40:01.0
U7879.001.DQD014E-P1-C6  PCI-X capable, 64 bit, 133MHz slot  0001:58:01.0
0001:58:01.1
```

List all empty PCI Hotplug slots: Type the following command to list all empty PCI Hotplug slots:

```
# lsslot -c pci -o
```

The resulting output looks similar to the following:

```
# Slot                     Description                         Device(s)
U7879.001.DQD014E-P1-C1    PCI-X capable, 64 bit, 133MHz slot  Empty
U7879.001.DQD014E-P1-C4    PCI-X capable, 64 bit, 133MHz slot  Empty
U7879.001.DQD014E-P1-C5    PCI-X capable, 64 bit, 133MHz slot  Empty
```

Show detailed information about a particular device: Type the following command to show detailed information about a particular device:

```
# lspci | grep xxxx:yy:zz.t
```

where *xxxx:yy:zz.t* is the number of the PCI Hotplug device. The resulting output looks similar to the following:

```
0001:40:01.0 Ethernet controller: Intel Corp. 82545EM Gigabit
Ethernet Controller (Copper) (rev 01)
```

Replacing a PCI Hotplug HBA

Syntax for the drslot_chrp_pci command: Use the **drslot_chrp_pci** tool as follows:

- Syntax: **drslot_chrp_pci -c slot {-r | -a | -R | -i } -s** *slot-name*
- Options:

  -r   Remove

  -a   Add

  -R   Replace

  -i   Identify

Replacing an HBA: Complete the following procedures to replace a PCI Hotplug HBA by using the drslot_chrp_pci command.

Note: In these procedures, the variable *slot-name* refers to the slot that contains the HBA that you are replacing.

1. Identify the PCI Hotplug slot: Type the following command to identify the PCI Hotplug slot:

```
# drslot_chrp_pci -i -s  slot-name
```

where *slot-name* is the name of the slot for the HBA you are replacing. (Example: U7879.001.DQD014E-P1-C3)

The light at slot *slot-name* begins flashing, and this message displays:

```
The visual indicator for the specified
PCI slot has been set to the identify
state. Press Enter to continue or
enter x to exit.
```

2. Hot unplug the HBA from the slot: Complete the following steps to hot unplug (remove) the HBA:

1. Type the following command:

```
# drslot_chrp_pci -r -s   slot-name
```

    This message displays:

```
The visual indicator for the specified
PCI slot has been set to the identify
state. Press Enter to continue or
enter x to exit.
```

2. Press **Enter**. This message displays:

```
The visual indicator for the specified
PCI slot has been set to the action state.
Remove the PCI card from the identified slot
and press Enter to continue.
```

3. Press **Enter**.
4. Physically remove the HBA from the slot.
5. Type the following command to verify that the slot is empty:

```
# lsslot -c pci -s    slot-name
```

    If the slot is empty, the resulting output looks similar to the following:

```
# Slot                    Description                      Device(s)
U7879.001.DQD014E-P1-C3   PCI-X capable, 64 bit, 133MHz slot   Empty
```

3. Hot plug the HBA into the slot: Complete the following steps to hot plug the HBA into the slot.

1. Type the following command:

```
# drslot_chrp_pci -a -s     slot-name
```

    This message displays:

```
The visual indicator for the specified
PCI slot has been set to the identify
state. Press Enter to continue or
enter x to exit.
```

2. Press **Enter**. This message displays:

```
The visual indicator for the specified
PCI slot has been set to the action state.
Insert the PCI card into the identified slot,

connect any devices to be configured
and press Enter to continue. Enter x to exit.
```

3. Insert the new HBA into the slot.

4. Type the following command to verify that the slot is no longer empty:

```
# lsslot -c pci -s      slot-name
```

If the slot is not empty, the resulting output looks similar to the following:

```
# Slot                  Description                              Device(s)
U7879.001.DQD014E-P1-C3  PCI-X capable, 64 bit, 133MHz slot      0001:40:01.0
```

Reassigning the WWPN on the DS4000 storage subsystem

After you replace the HBA, you must reassign its new WWPN on the DS4000 storage system. Complete the following steps:

1. If RDAC is installed, type the following command to recognize the new HBA:

```
# mppBusRescan
```

2. Start Storage Manager.
3. In the Mappings View of the Storage Manager Subsystem Management window, select the HBA Host Port that was previously assigned to the HBA that you replaced.
4. Right-click the HBA Host Port, then select **Replace...**. A new Storage Manager window opens.
5. In the new window, select the new host WWPN from the pull-down menu, or type the WWPN in the New Identifier field.
6. Click **OK** to close the window.
7. Type **mppBusRescan** again. Verify that the newly-replaced HBA recognizes the LUNs by checking the **lsscsi** output.

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 183
- "Renaming storage subsystems" on page 184
- "Creating a storage subsystem profile" on page 184
- "Performing other storage subsystem management tasks" on page 184

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 44.) 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 13.
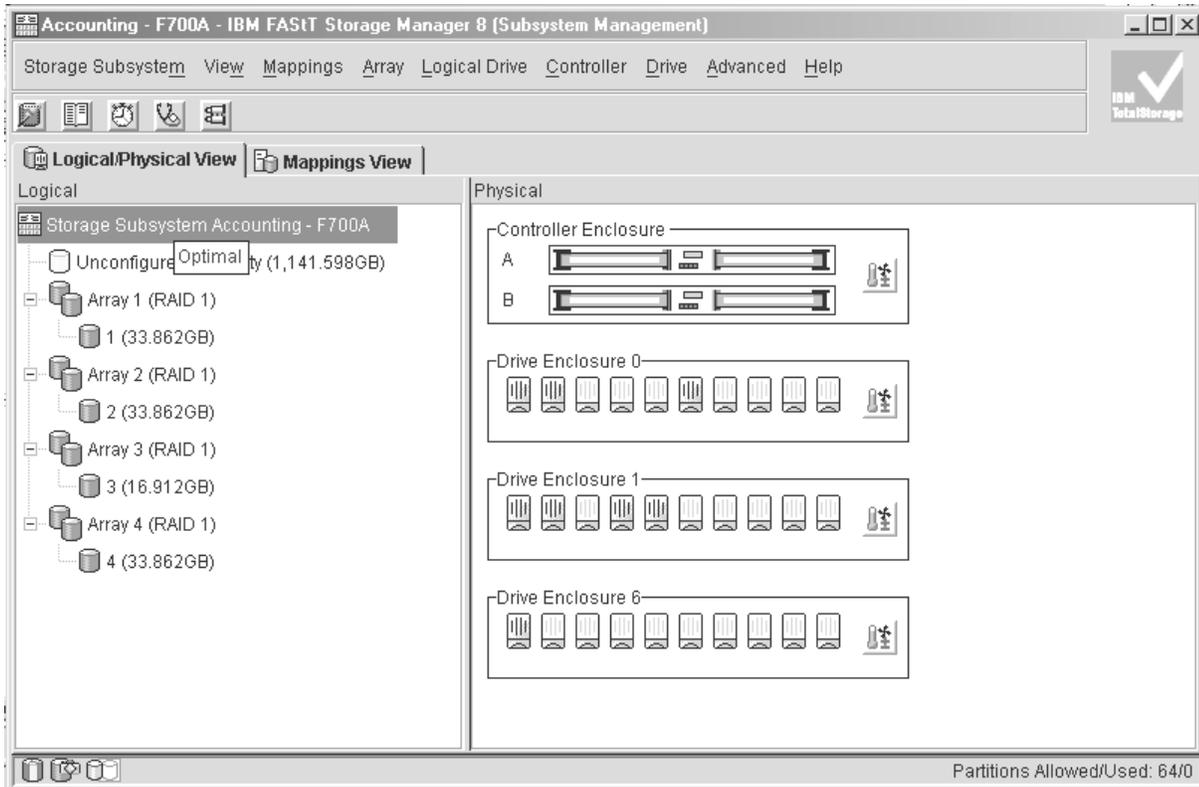
SJ001148

Figure 13. 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 storage subsystem is on the local subnetwork. If it is not, click **Edit —> Add Storage Subsystem** to add it. See the Enterprise Management window help topic on adding devices.

   > **Note:** Failure to add both controllers results in a partially-managed subsystem.

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...** in the Enterprise Management window. Click **Yes** when the dialog window appears.

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 183.

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 Logical (left) panel of the Subsystem Management window, highlight either **Free Capacity** or **Unconfigured Capacity**.

2. 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.
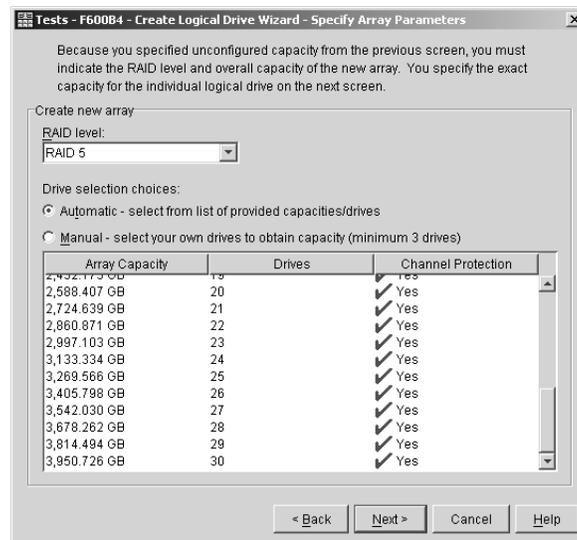
Figure 14. Create Logical Drive — Introduction window

3. 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.

4. 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.

5. In the Specify Capacity/Name (Create Logical Drive) window, specify the following logical drive parameters for the logical drive you want to create:

New logical drive capacity
> The new logical drive capacity can be either the entire unconfigured capacity in an array or a portion of the capacity.

Units You can choose from GB, MB, and TB.

Name The logical drive name must be unique in the storage subsystem.

Advanced logical drive parameters
> You can create the logical drive either using the DS4000 Storage Subsystem 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 Subsystem default settings, click the check box next to ″Use recommended settings,″ click **Next** then go to step 7. 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 6.

6. 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**.

7. 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:

Default mapping
> The Default mapping 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 Mappings view). If the Storage Partition feature is not enabled, you must specify the Default mapping 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.

8. 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.

9. If you do not want to create another logical drive, go to step 11. 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 10. If you do not want to create another logical drive, click **No** and continue with step11.

10. 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 5 on page 182.

11. The Completed (Create Logical Drive) window opens. Click **OK**.

12. 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 subsystem. 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 subsystem 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 the following sections for installation instructions:

- Chapter 2, "Installing and configuring storage management software on AIX host systems," on page 27
- Chapter 3, "Installing and configuring storage management software on HP-UX host systems," on page 77
- Chapter 4, "Installing and configuring storage management software on Solaris host systems," on page 99
- Chapter 5, "Installing and configuring storage management software on POWER-based Linux host systems," on page 133

Chapter 7. Using DS4000 Storage Manager with high-availability cluster services

The high-availability clustering services provided by DS4000 Storage Manager allow application services to continue when a hardware or software failure occurs. This system protects you from software failures as well as from the failure of a CPU, disk, or LAN component. If a component fails, its redundant partner component takes over cluster services and coordinates the transfer between components.

This chapter provides the following information:

- "Using cluster services on HP-UX systems" on page 188
- "Using cluster services on Solaris systems" on page 188
- "Using cluster services on AIX systems" on page 190

General information

You can use DS4000 Storage Manager with the following cluster services.

> **Important**
>
> The information in this document might not show up-to-date cluster software version levels.
>
> For the latest requirements and user information about using DS4000 Storage Manager with cluster services, see the readme file that is located in the DS4000 Installation CD for your host operating system. The most recent copies of the readme files are maintained at the following Web site:
>
> www-1.ibm.com/servers/storage/support/disk/
>
> Click the link for your storage subsystem. When the page opens, click the **Download** tab, then click the Storage Manager link, then click the **Storage Mgr** tab.
>
> You can also find more information on the DS4000 Interoperability matrix, which is maintained at the following Web site:
>
> www-1.ibm.com/servers/storage/disk/ds4000/interop-matrix.html

Operating System	High-Availability Cluster Service
AIX	HACMP 5.1, 5.2 PSSP 3.4 with GPFS 1.5 PSSP 3.5 with GPFS 2.1
HP-UX	HP MC/Service Guard Version 11.14
Solaris 8	VERITAS Cluster Server 3.5 or higher VERITAS Volume Manager 3.1, 3.2, 3.5 or higher VERITAS File System 3.5

Note: Clustering is not currently supported on Linux on POWER hosts with DS4000.

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

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

Using cluster services on HP-UX systems

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

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

See Appendix A, "MC/Service Guard configuration details," on page 199 for more information.

Using cluster services on Solaris systems

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

Important

The information in this document might not show up-to-date cluster software version levels. Check the Storage Manage readme file for Solaris at the following Web site for up-to-date information about clustering requirements:

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

Click the link for your storage subsystem. When the page opens, click the **Download** tab, then click the Storage Manager link, then click the **Storage Mgr** tab.

You can also find more information on the DS4000 Interoperability matrix, which is maintained at the following Web site:

www-1.ibm.com/servers/storage/disk/ds4000/interop-matrix.html

---

General hardware requirements

Each Solaris system in the cluster requires the following hardware:
- At least three Ethernet ports:
  – Two for the private network connections
  – At least one for the public network connection
- Two Fibre Channel host bus adapters for connection to the storage subsystem
- A SCSI connection for operating system disks
- Each VERITAS Cluster Server system requires at least 128 MB of RAM and 35 MB of free disk space

System dependencies

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

VERITAS Volume Manager and Data Multipathing

If you are using RDAC as your failover driver, and you are using the VERITAS Volume Manager, you must disable Data Multipathing (DMP). Complete one of the following procedures, according to your version of VERITAS Volume Manager.

Note: If you are using the VERITAS Cluster Server, you must change the stack size parameters. Open the /etc/system file in the Vi text editor and make the following changes:

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

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

Complete the following steps:

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

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

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

VERITAS Volume Manager 3.2 or later: For VERITAS Volume Manager 3.2 or later, you must disable DMP for DS4000 storage subsystems only.

Use option 17 and then option 5 of the submenu of the vxdiskadm utility to disable DMP for DS4000 storage subsystems. Other devices can use DMP.

RDAC IDs

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

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

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

2. Modify the `Rdac_HotAddIDs` line as follows:

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

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

Single points of failure

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

Using cluster services on AIX systems

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

> **Important**
>
> The information in this document might not show up-to-date cluster software version levels. Check the Storage Manage readme file for AIX at the following Web site for up-to-date information about clustering requirements:
>
> www-1.ibm.com/servers/storage/support/disk/
>
> Click the link for your storage subsystem. When the page opens, click the **Download** tab, then click the Storage Manager link, then click the **Storage Mgr** tab.
>
> You can also refer to the following Web sites for the most current information about AIX and clustering:
>
> ftp://ftp.software.ibm.com/storage/fastt/fastt500/HACMP_config_info.pdf
>
> www-1.ibm.com/servers/storage/disk/ds4000/interop-matrix.html
>
> http://publib.boulder.ibm.com/infocenter/clresctr/index.jsp

General hardware requirements

For hardware, software, firmware and RDAC requirements, see "AIX system requirements" on page 29.

High Availability Cluster Multi-Processing

This section contains general requirements and usage notes for High Availability Cluster Multi-Processing (HACMP) support with DS4000 Storage Manager.

Software requirements

The following HACMP software requirements are in addition to the requirements that are specified in "AIX system requirements" on page 29:

Table 30. HACMP software requirements

AIX versions	HACMP version	Required HACMP APAR
AIX 5.1, 5.2, 5.3	HACMP 5.1	IY53044
AIX 5.1, 5.2, 5.3	HACMP 5.2	IY58496

Configuration limitations

The following limitations apply to HACMP configurations:

- Switched fabric connection only; no direct connection is allowed between the host nodes and DS4000.
- HACMP C-SPOC cannot be used to add a DS4000 disk to AIX using the *Add a Disk to the Cluster* facility.
- HACMP C-SPOC does not support enhanced concurrent mode volume groups.
- Single node quorum is not supported in a two-node GPFS cluster with DS4000 disks in the configuration.

Other HACMP usage notes

The following notations are specific to HACMP environments:

- HACMP clusters can support from two to 32 servers on each DS4000 partition. If you run this kind of environment, be sure to read and understand the AIX device drivers queue depth settings that are described in Chapter 2, "Installing and configuring storage management software on AIX host systems," on page 27.
- You can attach non-clustered AIX hosts to a DS4000 that is running DS4000 Storage Manager and is attached to an HACMP 4.5 cluster. However, you must configure the non-clustered AIX hosts on separate host partitions on the DS4000.

Parallel System Support Programs and General Parallel File System

This section contains general requirements and usage notes for Parallel System Support Programs (PSSP) and General Parallel File System (GPFS) support with DS4000 Storage Manager.

Software requirements

The following combinations of PSSP and GPFS are supported:

Table 31. PSSP and GPFS software requirements

PSSP and GPFS versions	Required PSSP APARs
PSSP 3.5 with GPFS 2.1	**PSSP 3.5:** IY49062

Configuration limitations

The following limitations apply to PSSP and GPFS configurations:

- Direct connection is not allowed between the host node and a DS4000 storage subsystem. Only switched fabric connection is allowed.
- RVSD clusters can support up to two IBM Virtual Shared Disk and RVSD servers for each DS4000 partition.
- Single node quorum is not supported in a dual-node GPFS cluster with DS4000 disks in the configuration.
- Heterogeneous configurations are not supported.

Other PSSP and GPFS usage notes

In GPFS file systems, the following DS4000 cache settings are supported:

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

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

GPFS, PSSP, and HACMP cluster configuration diagrams

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

A preferred path to a volume is determined when the volume is created and distributed across a DS4000 controller. The controller to which it is assigned determines which path is preferred or active for I/O transfer. Volumes can, and in most cases should, be assigned to both controllers, balancing the I/O load across HBAs and DS4000 controllers.

Figure 15 shows a cluster configuration that contains a single DS4000 storage subsystem, with one to four partitions.
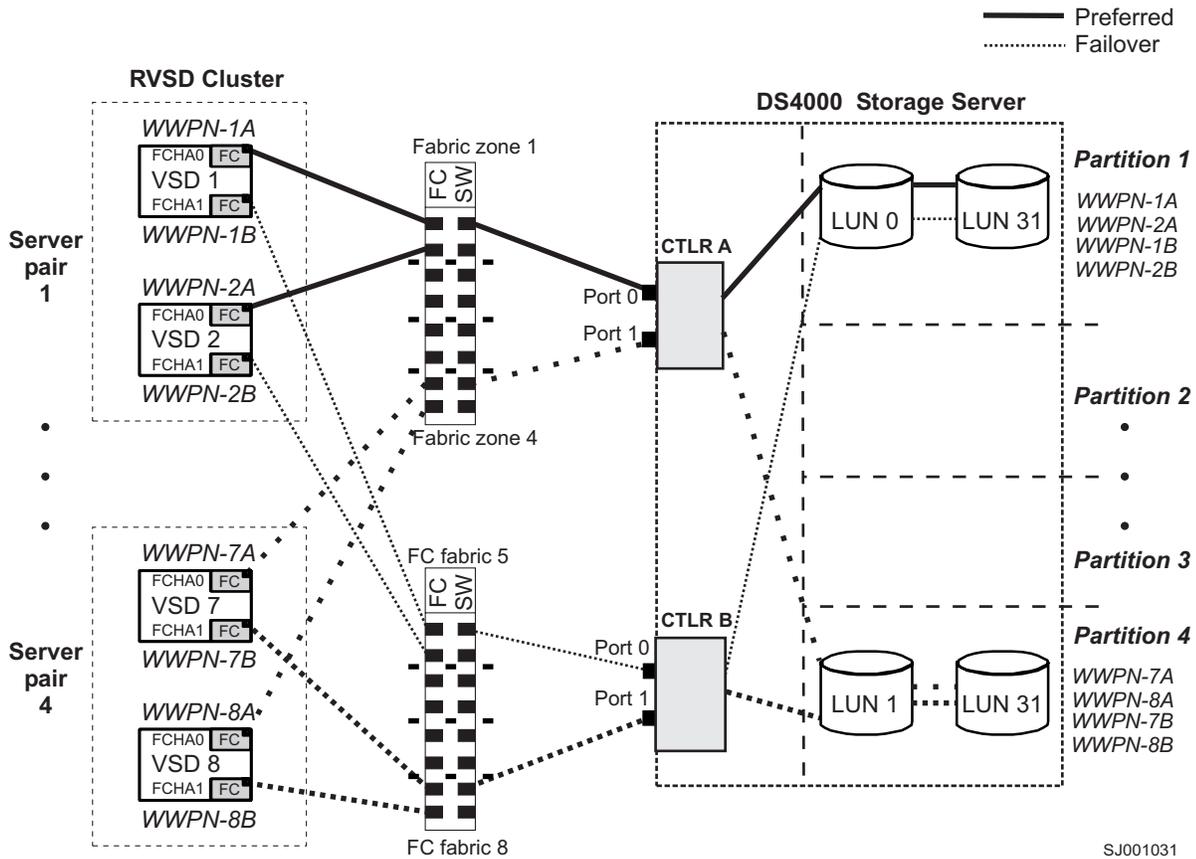
Figure 15. Cluster configuration with single DS4000 storage subsystem — one to four partitions

Figure 16 on page 193 shows a cluster configuration that contains three DS4000 storage subsystems, with one partition on each storage subsystem.
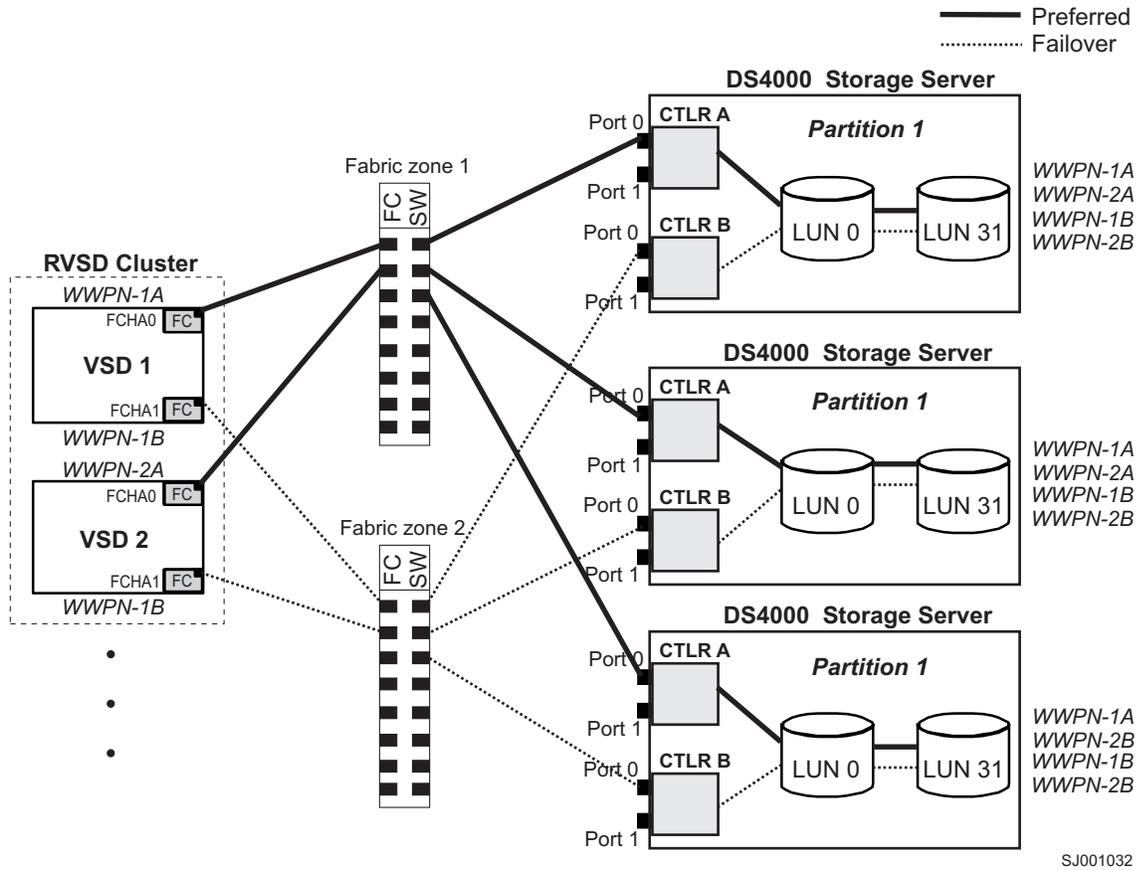
Figure 16. Cluster configuration with three DS4000 storage subsystems — one partition per DS4000

Figure 17 on page 194 shows a cluster configuration that contains four DS4000 storage subsystems, with one partition on each storage subsystem.
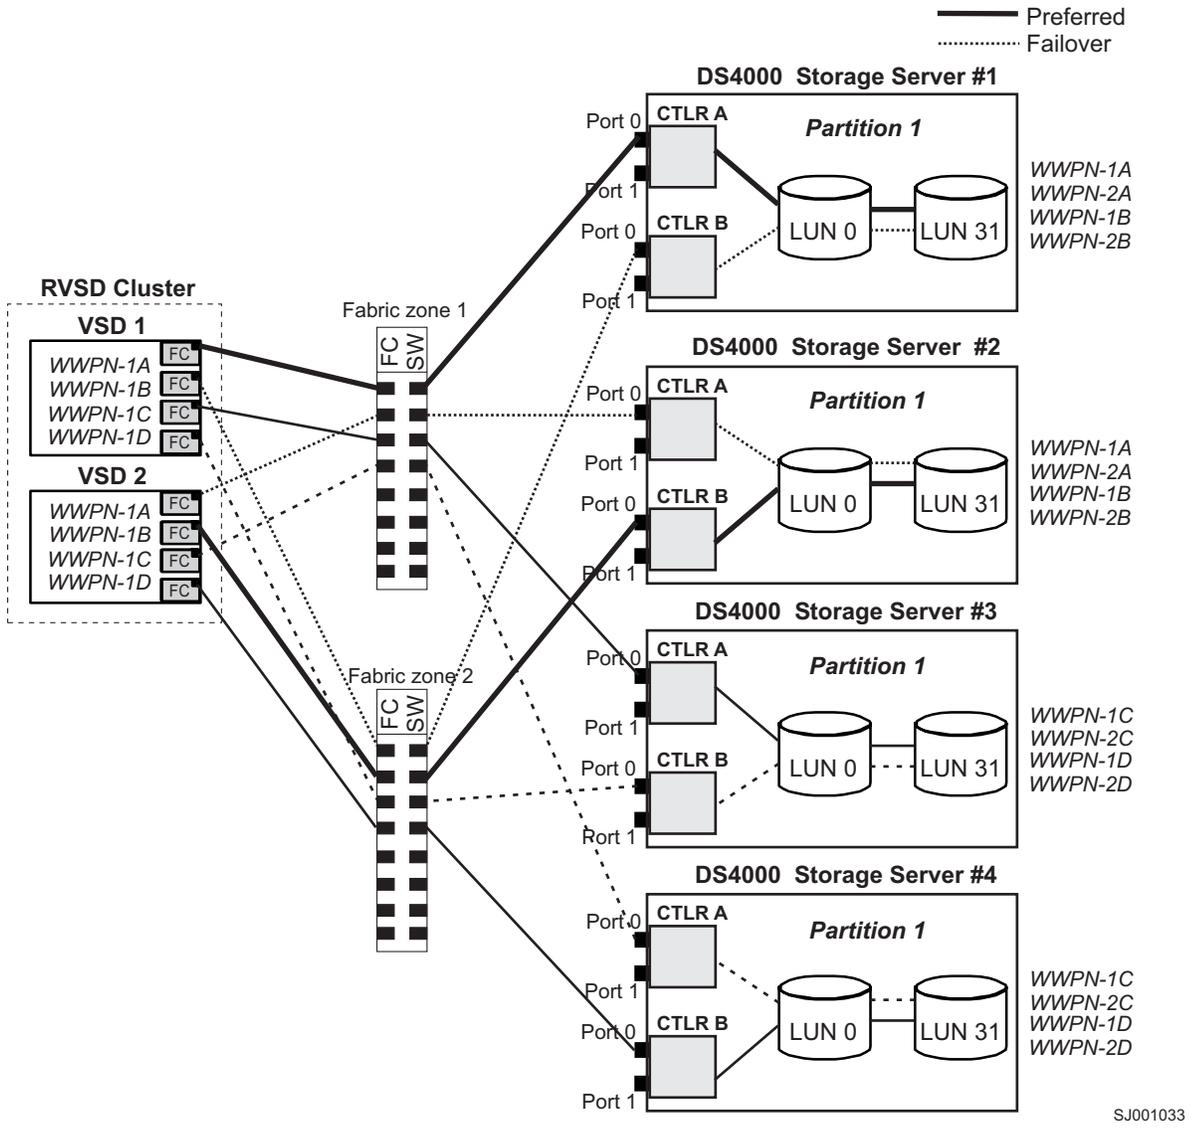
Figure 17. Cluster configuration with four DS4000 storage subsystems — one partition per DS4000

Figure 18 on page 195 shows a cluster configuration that contains two DS4000 storage subsystems, with two partitions on each storage subsystem.

DS4000 Storage Server #1

Figure 18. RVSD cluster configuration with two DS4000 storage subsystems — two partitions per DS4000

Figure 19 on page 196 shows an HACMP/GPFS cluster configuration that contains a single DS4000 storage subsystem, with one partition.
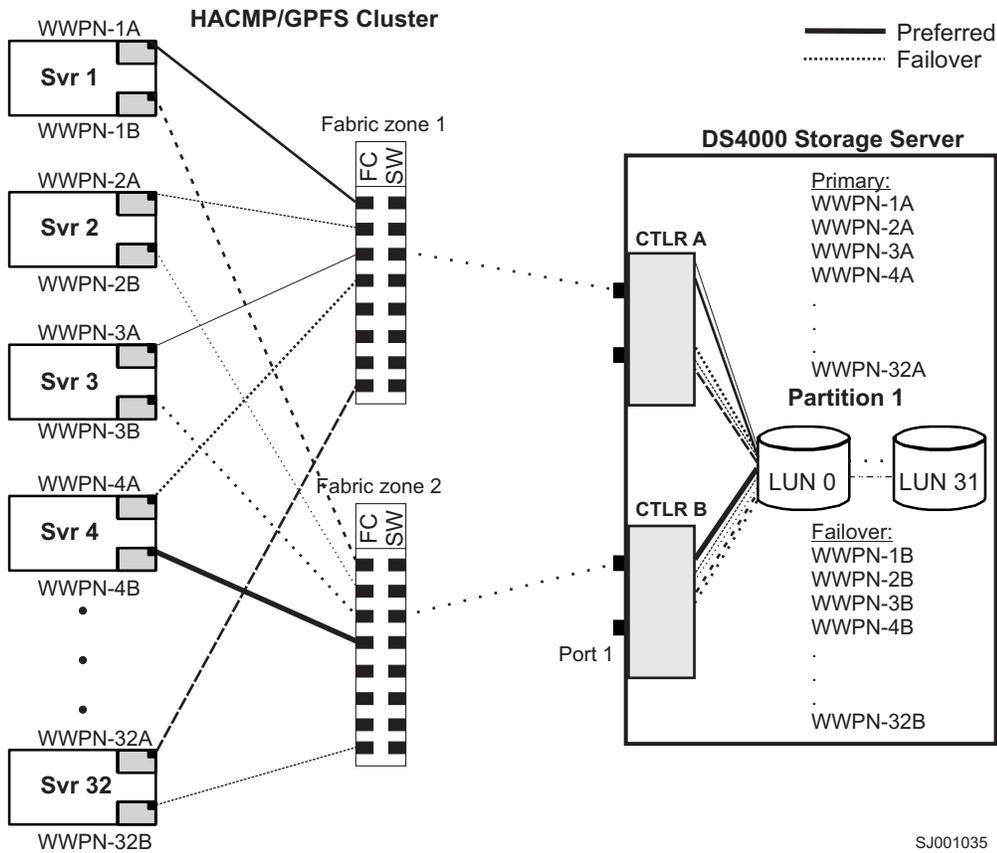
Figure 19. HACMP/GPFS cluster configuration with one DS4000 storage subsystem — one partition

Figure 20 on page 197 shows an HACMP/GPFS cluster configuration that contains two DS4000 storage subsystems, with two partitions on each storage subsystem.
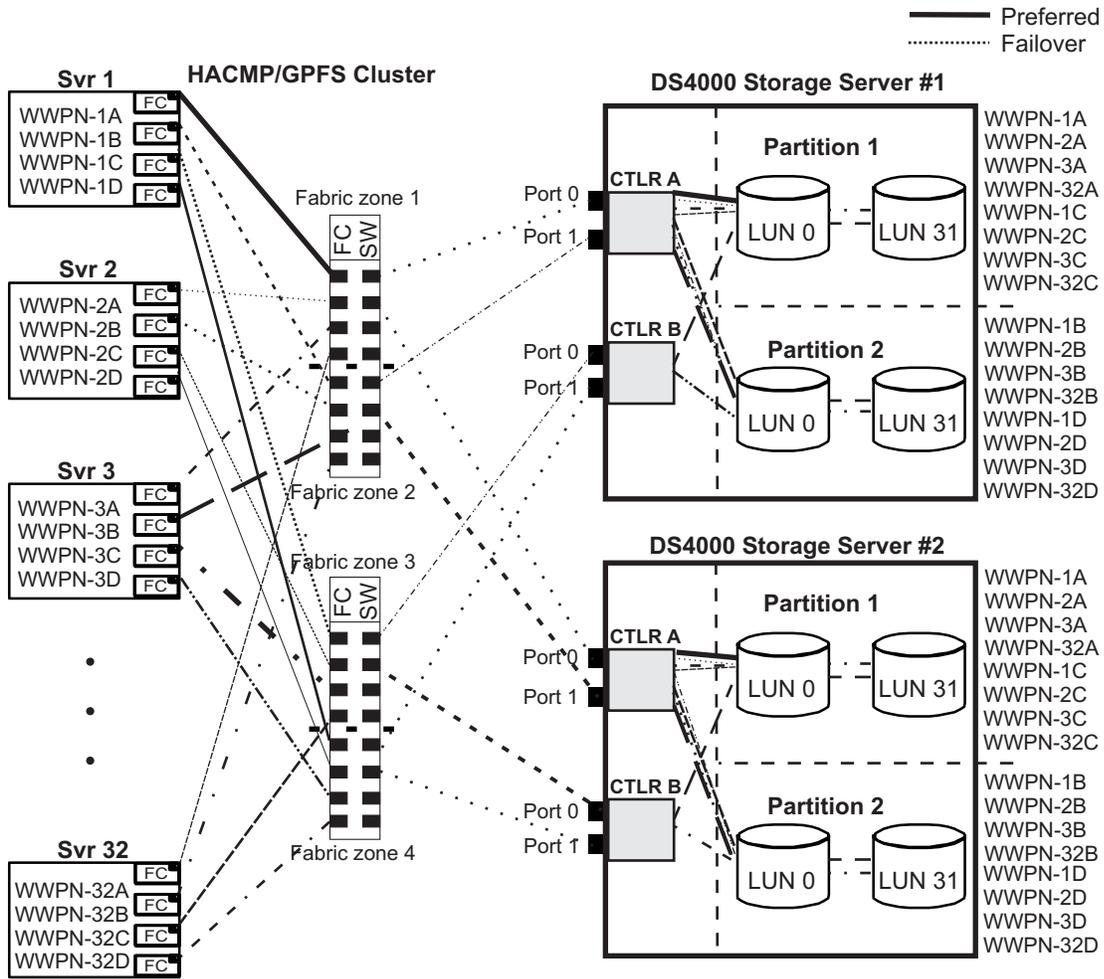
Figure 20. HACMP/GPFS cluster configuration with two DS4000 storage subsystems — two partitions per DS4000

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 32. 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/genjniconf* reconfigure script from the Solaris shell.

```
# /etc/raid/bin/genjniconf
```

Table 33. 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 33. 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/genjniconf* reconfigure script from the Solaris shell.

```
# /etc/raid/bin/genjniconf
```

Table 34. 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/genjniconf* reconfigure script from the Solaris shell.

```
# /etc/raid/bin/genjniconf
```

- Set *portEnabled = 1;* only when you see JNI cards entering non-participating mode in the */var/adm/messages* file. Under that condition:
  1. Set *FcPortCfgEnabled = 1;*

2. Restart the host.
3. Set *FcPortCfgEnabled = 0;*
4. Restart the host again.

When you have done so, check */var/adm/messages* to be sure that it sets the JNI cards to Fabric or Loop mode.

Table 35. Configuration settings for FC64-1063

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

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

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

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 36.

Note: In Table 36, 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 36.

Table 36. 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 ESM firmware to an DS4000 EXP storage expansion enclosure and current drive firmware to a DS4000 storage subsystem, using the following procedures:

- "Verifying and downloading ESM firmware"
- "Downloading drive firmware" on page 208

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 206

Important: If you are using DS4000 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 follow the links to the readmes.

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, click **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 drive enclosure.

     Tip: Each 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 DS4000 readme file, which is located at the following Web site:

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

Click the link for your storage subsystem. When the page opens, click the Download tab, then follow the links to the readmes.

Complete the following procedures to download ESM firmware:
- "Preparing for the download"
- "Downloading the ESM firmware" on page 207

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 207.

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 subsystem.

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 Subsystem profiles in the same location as the logical drives that are defined on your DS4000 Storage Subsystem. 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 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 DS4000 readme file, which is located at the following Web site:

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

Click the link for your storage subsystem. When the page opens, click the Download tab, then follow the links to the readmes.

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 enclosure.

When the status field of each 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 Subsystem.
   - 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 drive firmware with Storage Manager version 9.1x

To download drive firmware for using DS4000 Storage Manager version 9.1x, 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 9.1x 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 subsystem 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 is 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 on page 208 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 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 15.
         1) Click the **Select all** button.
         2) Click **Save the Storage Subsystem Event Log**.

   When the **Close** button appears active, the drive firmware download process is complete.

15. Click **Close** to exit the Download Progress window.

16. Use either of the following procedures to determine or verify what level of drive firmware resides on a particular drive:
   • Right-click 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 **Storage Subsystem —> View Profile** in the Logical/Physical View of the Subsystem Management window.

Downloading drive firmware with Storage Manager version 8.x

To download drive firmware for using 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. Limitations

This sections in this appendix contain descriptions of potential problems and ways to work around them.

AIX

The following information might apply to your storage subsystem.

1. Storage Manager prompts you for a password if all hard disk drives are removed from the storage subsystem. After removing all drives from a storage subsystem (or not powering on drives and only the controller), the IBM DS4000 Storage Manager software will prompt you for a password when trying to start the software or trying to perform protected operations. Any password entered will fail.

   Password information is stored on a reserved area of each hard disk drive on the storage subsystem. Each drive stores a mirrored copy of the password data. When there are no drives in the storage subsystem, the storage management software does not find the password data when you attempt password protected operations. Add one drive to the storage subsystem and retry the operation.

2. Auto LUN/Array Transfer is not supported. If a controller fails over arrays to the alternate controller and the failed controller is replaced and brought back online, the arrays do not automatically transfer to the preferred controller. This must be done manually by redistributing arrays.

3. If the controller in the IBM DS4000 Storage Server with firmware 05.4x.xx.xx is connected to a SAN with more than 256 host connections per controller port, that can log into the controllers, one must either use the switch zoning function or rearrange the SAN connections so that only the maximum number of host connections can be seen and log into the controller.

4. Dynamically changing the cache read ahead multiplier setting via the storage manager client may result in an error message indicating the operation has failed. Despite the error message, the request to change the cache read ahead multiplier will remain queued to the controllers and will eventually complete when I/O activity to the storage subsystem has subsided. In order to avoid the error message and any delays in the cache read ahead multiplier value changing, it is recommended to change the cache read ahead multiplier when the subsystem is in a quiesced state.

HP-UX

The following information might apply to your storage subsystem.

1. If you are upgrading controller firmware in IBM FAStT500 storage controller with 512MB Cache, you will have to stop I/O and reboot both controllers before upgrading firmware followed by NVSRAM. You can reboot the controllers either by power cycling the storage subsystem or by pressing the Reset button on both controllers. The Reset button is located on the right of the indicator lights in the opening on the front of each controller unit.

   To determine Storage Controller cache level, perform the following:

   a. On a storage management station start the client software and select a storage subsystem.

   b. Select Tools -> Manage Device to open the Subsystem Management Window

c. Highlight the Storage Subsystem that you are upgrading and then select Storage Subsystem -> Profile. It may take a little time for it to create the Storage Subsystem profile.

d. Under the Controllers heading is an entry that displays information similar to the following example:

```
Cache/processor size (MB):  256/32
```

If this setting is listed as 512/32 you will need to arrange to perform this upgrade procedure at a time that you can stop I/O to the Storage Subsystem in order to ensure a successful firmware download and installation.

2. Storage Manager prompts you for a password if all hard disk drives are removed from the storage subsystem. After removing all drives from a storage subsystem (or not powering on drives and only the controller), the IBM DS4000 Storage Manager software will prompt you for a password when trying to start the software or trying to perform protected operations. Any password entered will fail.

Password information is stored on a reserved area of each hard disk drive on the storage subsystem. Each drive stores a mirrored copy of the password data. When there are no drives in the storage subsystem, the storage management software does not find the password data when you attempt password protected operations. Add one drive to the storage subsystem and retry the operation.

Since the RAID controller unit should not be powered on until the expansion units are attached and powered on, this should be done by powering off the RAID controller unit, properly attaching all expansion units (with drives installed), to the RAID controller unit and powering them on, and then powering the RAID controller unit back on. At this point the storage subsystems can be rediscovered and managed using the IBM DS4000 Storage Manager 8.3x software.

3. Display issues when using 256 color settings. A brownish hash pattern is displayed on the windows and in online help when running in 256 color mode. Run the Storage Manager 8 application in a higher display mode.

4. If you get Error 21 - The Operation cannot complete because a ... is in progress error when you issue a Subsystem Management command from the IBM DS4000 Storage Manager Client, the command may be cached. If the command is cached, the IBM DS4000 Storage Manager Client will run the command as soon as the controller resources are available. Wait for at least ten minutes and check to see whether the command has been run. If not, re-attempt the command.

5. The IBM FAStT500, DS4400 or DS4500 Storage Server hardware and software is designed to have redundant paths from the host system to ensure the availability of storage. However, it is possible to connect a single host adapter from the host system to one of the controller minihubs and then connect this minihub to the minihub of the other controller.

This configuration will create a FC-AL loop between the host adapter and the two IBM FAStT500, DS4400 or DS4500 Storage Server controllers and no redundant path to the host system. When you connect the two minihubs together, always connect the IN port of one minihub to the OUT port of the other minihub or vice versa. Do not connect the IN port of one to the IN port of the other minihub or the OUT port of one minihub to the OUT port of the other minihub.

6. DS4000 now supports 127 LUNs on HP-UX. However, the number of actual LUNs might vary from HP-UX host to host. The number of LUNs supported varies depending on the number of processors and memory.

7. Dynamically changing the cache read ahead multiplier setting via the storage manager client may result in an error message indicating the operation has failed. Despite the error message, the request to change the cache read ahead multiplier will remain queued to the controllers and will eventually complete when I/O activity to the storage subsystem has subsided. In order to avoid the error message and any delays in the cache read ahead multiplier value changing, it is recommended to change the cache read ahead multiplier when the subsystem is in a quiesced state.

Solaris

1. Upgrading FAStT500 with 512MB Cache. If you are upgrading controller firmware in IBM FAStT500 storage controller (MT 3552), you will have to stop I/O and reboot both controllers before upgrading NVSRAM followed by Firmware. You can reboot the controller by either power cycle the storage subsystem or by pressing the Reset button on both controllers. The Reset button is located on the right of the indicator lights in the opening on the front of each controller unit.

   To determine Storage Controller cache level, perform the following:

   a. On a storage management station start the client software and select a storage subsystem.

   b. Select Tools -> Manage Device to open the Subsystem Management Window

   c. Highlight the Storage Subsystem that you are upgrading and then select Storage Subsystem -> Profile. It may take a little time for it to create the Storage Subsystem profile.

   d. Under the Controllers heading is an entry that displays information similar to the following example:

   ```
   Cache/processor size (MB):  256/32
   ```

   If this setting is listed as 512/32 you will need to arrange to perform this upgrade procedure at a time that you can stop I/O to the Storage Subsystem in order to ensure a successful firmware download and installation.

2. Storage Manager prompts you for a password if all hard disk drives are removed from the storage subsystem. After removing all drives from a storage subsystem (or not powering on drives and only the controller), the IBM DS4000 Storage Manager software will prompt you for a password when trying to start the software or trying to perform protected operations. Any password entered will fail.

   Password information is stored on a reserved area of each hard disk drive on the storage subsystem. Each drive stores a mirrored copy of the password data. When there are no drives in the storage subsystem, the storage management software does not find the password data when you attempt password protected operations. Add one drive to the storage subsystem and retry the operation.

   Since the RAID controller unit should not be powered on until the expansion units are attached and powered on, this should be done by powering off the RAID controller unit, properly attaching all expansion units (with drives installed), to the RAID controller unit and powering them on, and then powering the RAID

controller unit back on. At this point the storage subsystems can be rediscovered and managed using the IBM DS4000 Storage Manager 8.3x software.

3. Auto LUN/Array Transfer is not supported. If a controller fails over arrays to the alternate controller and the failed controller is replaced and brought back online, the arrays do not automatically transfer to the preferred controller. This must be done manually by redistributing arrays.

4. Display issues when using 256 color settings. A brownish hash pattern is displayed on the windows and in online help when running in 256 color mode. Run the Storage Manager 8 application in a higher display mode.

5. If you get Error 21 - The Operation cannot complete because a ... is in progress error when you issue a Subsystem Management command from the IBM DS4000 Storage Manager Client, the command may be cached. If the command is cached, the IBM DS4000 Storage Manager Client will run the command as soon as the controller resources are available. Wait for at least ten minutes and check to see whether the command has been run. If not, re-attempt the command.

6. The IBM FAStT500, DS4400 or DS4500 Storage Server hardware and software is designed to have redundant paths from the host system to ensure the availability of storage. However, it is possible to connect a single host adapter from the host system to one of the controller minihubs and then connect this minihub to the minihub of the other controller.

   This configuration will create a FC-AL loop between the host adapter and the two IBM FAStT500, DS4400 or DS4500 Storage Server controllers and no redundant path to the host system. When you connect the two minihubs together, always connect the IN port of one minihub to the OUT port of the other minihub or vice versa. Do not connect the IN port of one to the IN port of the other minihub or the OUT port of one minihub to the OUT port of the other minihub.

Appendix E. Additional DS4000 documentation

The following tables present an overview of the DS4500, DS4400, DS4300, DS4100, and DS4800 Storage Subsystem 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 37 associates each document in the DS4500 (previously FAStT900) Fibre Channel Storage Server library with its related common user tasks.

Table 37. 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**
IBM TotalStorage DS4500 Installation and Support Guide, GC26-7727	✔	✔		✔		
IBM TotalStorage DS4500 Fibre Channel Cabling Instructions, GC26-7729	✔	✔				
IBM TotalStorage DS4500 Storage Server User's Guide, GC26-7726				✔	✔	✔
IBM TotalStorage DS4500 Rack Mounting Instructions, GC26-7728	✔	✔				

DS4400 Fibre Channel Storage Server library

Table 38 associates each document in the DS4400 (previously FAStT700) Fibre Channel Storage Server library with its related common user tasks.

Table 38. 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**
IBM DS4400 Fibre Channel Storage Server User's Guide, GC26-7730	✔	✔		✔	✔	✔
IBM DS4400 Fibre Channel Storage Server Installation and Support Guide, GC26-7731	✔	✔		✔	✔	
IBM DS4400 Fibre Channel Cabling Instructions, GC26-7732	✔	✔				

DS4300 Fibre Channel Storage Server library

Table 39 associates each document in the DS4300 (previously FAStT600) Fibre Channel Storage Server library with its related common user tasks.

Table 39. 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**
IBM TotalStorage DS4300 Fibre Channel Storage Server Installation and User's Guide, GC26-7722	✔	✔		✔		
IBM TotalStorage DS4300 Rack Mounting Instructions, GC26-7724	✔	✔				
IBM TotalStorage DS4300 Fibre Channel Cabling Instructions, GC26-7725	✔	✔				
IBM TotalStorage DS4300 SCU Base Upgrade Kit, GC26-7740		✔	✔			
IBM TotalStorage DS4300 SCU Turbo Upgrade Kit, GC26-7741		✔	✔			
IBM TotalStorage DS4300 Turbo Models 6LU/6LX Upgrade Kit, GC26-7723		✔	✔			

DS4100 SATA Storage Server library

Table 40 associates each document in the DS4100 (previously FAStT100) SATA Storage Server library with its related common user tasks.

Table 40. 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**
IBM TotalStorage DS4100 Installation, User's and Maintenance Guide, GC26-7733	✔	✔		✔	✔	✔
IBM TotalStorage DS4100 Cabling Guide, 24P8973		✔				

DS4800 Storage Subsystem library

Table 41 associates each document in the DS4800 Storage Subsystem library with its related common user tasks.

Table 41. TotalStorage DS4800 Storage Subsystem document titles by user tasks

Title	User Tasks					
	Planning	**Hardware Installation**	**Software Installation**	**Configuration**	**Operation and Administration**	**Diagnosis and Maintenance**
IBM TotalStorage DS4800 Installation, User's and Maintenance Guide, GC26-7748	✔	✔		✔	✔	✔
IBM TotalStorage DS4800 Installation and Cabling Overview, 22R4738		✔				

Other DS4000 and DS4000-related documents

Table 42 associates each of the following documents with its related common user tasks.

Table 42. 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**
IBM Safety Information, 33P2406					✔	
IBM TotalStorage DS4000 Quick Start Guide, GC26-7738	✔	✔				
IBM TotalStorage DS4000 Hardware Maintenance Manual, GC26-7702						✔
IBM TotalStorage DS4000 Problem Determination Guide, GC26-7703						✔
IBM Fibre Channel Planning and Integration: User's Guide and Service Information, SC23-4329	✔	✔			✔	✔
IBM TotalStorage DS4000 FC2-133 Host Bus Adapter Installation and User's Guide, GC26-7736		✔			✔	
IBM TotalStorage DS4000 FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7737		✔			✔	
IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview GC26-7713	✔	✔	✔	✔		
Fibre Channel Solutions - IBM DS4000 EXP500 Installation and User's Guide, 59P5637	✔	✔		✔	✔	✔

Table 42. 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
IBM TotalStorage DS4000 EXP700 and EXP710 Storage Expansion Enclosures Installation, User's, and Maintenance Guide, GC26-7735	✔	✔		✔	✔	✔
IBM TotalStorage DS4000 Hard Drive and Storage Expansion Enclosures Installation and Migration Guide, GC26-7704	✔	✔				
IBM DS4000 Management Suite Java User's Guide, 32P0081					✔	✔
IBM Netfinity Fibre Channel Cabling Instructions, 19K0906		✔				
IBM Fibre Channel SAN Configuration Setup Guide, 25P2509	✔		✔	✔	✔	

Appendix F. 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's installation and user's guide.

Note: You can find links to switch documentation at the following Web site:

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

Figure 21 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 port through a fibre channel switch.
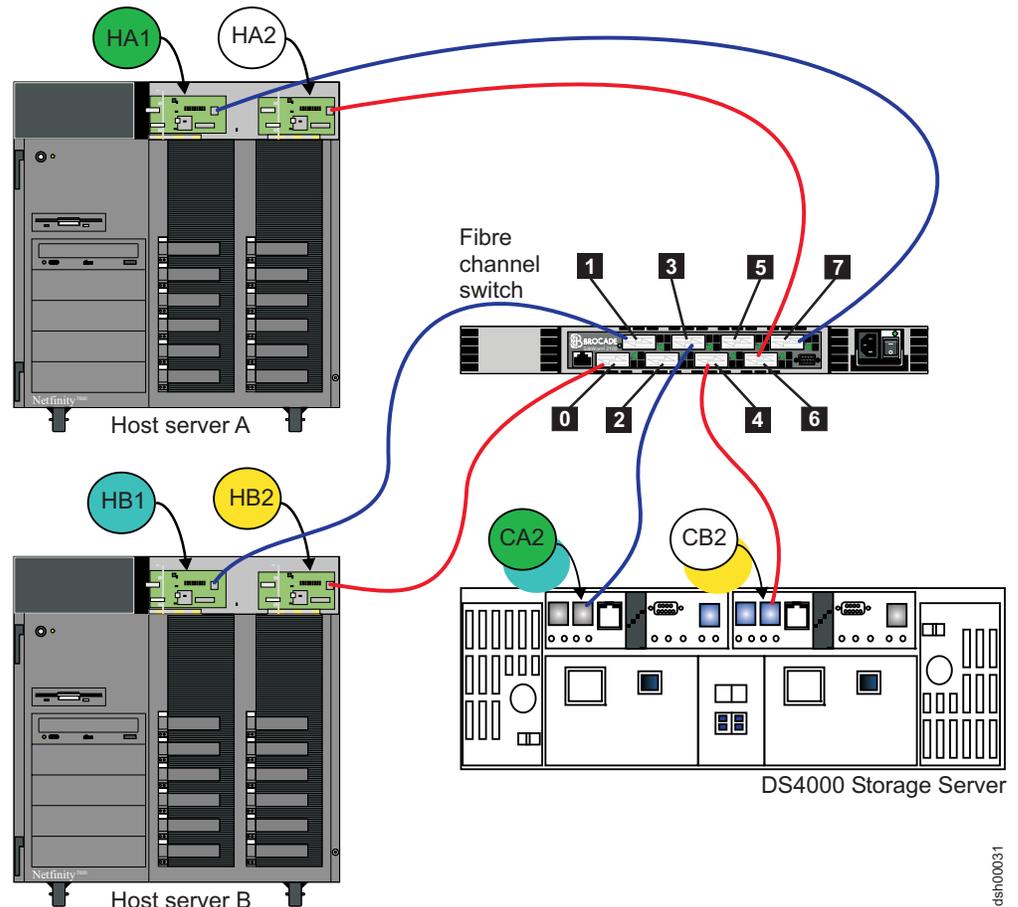
Figure 21. Fibre channel fabric switch environment connecting host servers to a DS4000 Storage Server

223

Legend:

Controller	Zone	Port	Connection to and from	Description
A	1	**7** **3**	HA1 CA2	This will zone the fibre channel HBA port HA1 in host server A to port 2 of the DS4000 Storage Server controller A.
	2	**1** **3**	HB1 CA2	This will zone the fibre channel HBA port HB1 in host server B to port 2 of the DS4000 Storage Server controller A.
B	3	**6** **4**	HA2 CB2	This will zone the fibre channel HBA port HA2 in host server A to port 2 of the DS4000 Storage Server controller B.
	4	**0** **4**	HB2 CB2	This will zone the fibre channel HBA port HB2 in host server B to port 2 of the DS4000 Storage Server controller B.

Appendix G. 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 43 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 43. 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 43. DS4000 Storage Manager alternate keyboard operations (continued)

Short cut	Action
Ctrl+Spacebar (Contiguous/Non-contiguous) AMW Logical/Physical View	Select multiple drives in the Physical View. 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. If you press Spacebar alone when multiple drives are selected then all selections are removed. Use the Ctrl+Spacebar combination to deselect a drive when multiple drives are selected. This behavior is the same for contiguous and non-contiguous selection of drives.
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*.

FAStT MSJ. See *FAStT Management Suite Java*.

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

FC. See *fibre channel*.

FC-AL. See *arbitrated loop*.

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

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

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

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

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

FlashCopy. A premium feature for 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 redundancy information.

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 know 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 packed-switched communication networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

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

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

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

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

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

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

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   xiii
access volume/logical drive   4, 13, 31, 58, 85, 108, 139
additional storage
   See storage expansion enclosures
address
   for e-mail comments   xix
   for reader comment form   xix
   of IBM   xix
   of the IBM director of licensing   227
agent software
   AIX   33
   HP-UX   84
   package info   10
   Red Hat Linux   149
   Solaris   106
AIX
   boot restrictions   31
   client software   32
   concurrent firmware download   35
   configuration of storage subsystems   34
   direct-attached configuration   41
   dynamic capacity expansion (DCE)   65
   dynamic volume expansion (DVE)   65
   fast I/O failure   63
   FCP disk array errors   72
   fget_config command   50
   firmware
      requirements   27
      upgrading   38
   hardware requirements   27, 29
   identifying
      controller ID numbers   46
      device names and bus numbers   48
      FlashCopy logical drive types   54
      logical drives by operating system device
         names   53
   logical drives, redistributing in case of failure   75
   lsattr command   52
   lsdev command   48
   NVSRAM
      requirements   27
      upgrading   38
   operating system requirements   27
   RDAC driver
      configuring   44
      installing   42, 44
      requirements   42
   requirements
      client software   32
      firmware   27
      hardware   27, 29
      operating system   27
      RDAC driver   42
      software   29
      system   29
   restrictions   30

AIX   *(continued)*
   SAN-attached configuration   41
   SMagent   33
   SMclient   33
   SMruntime   33
   software requirements   29
   staged controller firmware download   35
   system requirements   29
   verifying configuration   46
alert notifications, configuring   183
appware
   AIX firmware   27
   HP-UX firmware   77
   Red Hat Linux firmware   133
   Solaris firmware   99
   SUSE Linux firmware   133
attributes
   dac   61, 62
   dar   61
   hdisk   60, 61, 62
   LUN   60, 61, 62
   RDAC driver   44
audience of this document   xiii
Auto Volume Transfer (AVT)   97, 132

B

boot, AIX restrictions   31
BOOTP server
   sample network   15, 20
bus numbers and device names, identifying   48

C

cache mirroring   63, 74, 191
capture state information feature   184
catastrophic failure   184
check drive firmware level   209, 210
client software
   AIX   32, 33
   HP-UX   82, 83
   package info   9, 10
   Red Hat Linux   147, 148
      prerequisites   144
   Solaris   104, 105
cluster services
   AIX requirements   190, 191
   configurations   179
   HACMP ES and ESCRM   190
   hardware requirements   188
   HP-UX requirements   188
   MC/Service Guard   187, 199
   PSSP with GPFS   191
   Solaris requirements   188, 189
   system dependencies   187
   VERITAS Cluster Server   189
   VERITAS Volume Manager   189

239

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

IBM TotalStorage DS4000 Storage Manager Version 9
Installation and Support Guide for AIX, HP-UX, Solaris, and Linux on POWER

Publication No. GC26-7705-02

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

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.

IBM ®

Fold and Tape   **Please do not staple**   Fold and Tape
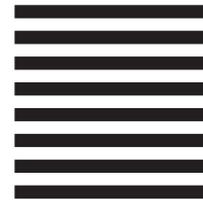
NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

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

IBM ®

Printed in USA

Spine information:

IBM TotalStorage DS4000 Storage
Manager Version 9

DS4000 Storage Manager Version 9 Installation and
Support Guide for AIX, HP-UX, Solaris, and Linux on
POWER