System Storage DS3000 Storage Manager Version 10

Installation and Support Guide for IBM AIX, Linux on POWER, and Sun Solaris

System Storage DS3000 Storage Manager Version 10

Installation and Support Guide for IBM AIX, Linux on POWER, and Sun Solaris

Note: Before using this information and the product it supports, read the general information in Appendix F, “Notices,” on page 129.

Fourth Edition (November 2008)

Contents

Figures

Tables

Chapter 1. Introduction

The IBM® System Storage™ DS3000 Storage Manager Version 10 (referred to throughout this document as *Storage Manager software*) is host software that you can use to manage your storage subsystems.

This document provides information about how to prepare, install, configure, and work with the Storage Manager software in each of the following operating systems:

- IBM AIX® version 5.2 and IBM AIX version 5.3, and IBM AIX version 6.1
- Red Hat Enterprise Linux® 4 (RHEL 4)
- Red Hat® Enterprise Linux 5 (RHEL 5)
- SUSE Linux Enterprise Server 9 (SLES 9)
- SUSE Linux Enterprise Server 10 (SLES 10) and SUSE Linux Enterprise Server 10 Service Pack 1
- Sun Solaris 10

Note: Storage Manager 10.35 does not contain a Solaris client that can be used to manage DS3000 storage subsystems. For information about how to set up the Solaris 10 Multiplexed I/O (MPxIO) failover driver, see Appendix C, “Solaris failover drivers,” on page 119.

For an interoperability matrix that lists the operating systems that are supported by the Storage Manager software, see the following Web sites:

- For the DS3200: http://www.ibm.com/systems/storage/disk/ds3000/ds3200/
- For the DS3300: http://www.ibm.com/systems/storage/disk/ds3000/ds3300/
- For the DS3400: http://www.ibm.com/systems/storage/disk/ds3000/ds3400/

Obtaining the documentation from the IBM Support Web site

If firmware and documentation updates are available, you can download them from the IBM support Web site. The Storage Manager software might have features that are not described in the documentation that comes with the unit, and the documentation might be updated occasionally to include information about those features, or technical updates might be available to provide additional information that is not included in the storage subsystem documentation. To check for updates, complete the following steps:

1. Go to http://www.ibm.com/servers/storage/support/.
2. Under **Select your product**, in the **Product Family** field, click **Disk systems**.
3. In the **Product** field, click **IBM System Storage DS3200**, **IBM System Storage DS3300**, or **IBM System Storage DS3400**, as applicable.
4. Click **Go**.
5. Make the following selections:
   - For firmware updates and readme files, click the **Download** tab.
   - For documentation updates, click the **Install and use** tab.

Note: Changes are made periodically to the IBM Web site. Procedures for locating firmware and documentation might vary slightly from what is described in this document.

Notices in this document

The following notices are used in this document:

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

Storage Manager software online help and diagnostics

When you are finished with the installation process, see the following online help systems:

- Enterprise Management window help

  Use this online help system to learn more about working with the entire management domain.
- Subsystem Management window help

  Use this online help system to learn more about managing individual storage subsystems.

You can access these help systems from within the Storage Manager Client. From either an Enterprise Management or Subsystem Management window, click **Help** or press F1.

Premium features

Premium features provide enhanced functionality that enables you to perform tasks that might be limited or unavailable with your base storage subsystem model. For more information, see Chapter 7, “Enabling and using premium features,” on page 109.

You can purchase the following premium features for a DS3000 storage subsystem:

- DS3000 FlashCopy® Expansion License
- DS3000 Volume Copy License
- DS3000 FlashCopy Volume Copy License
- DS3000 Partition Expansion License

FlashCopy

The FlashCopy feature supports the creation and management of FlashCopy logical drives. A FlashCopy logical drive is a logical point-in-time image of another logical drive, called a base logical drive, that is in the storage subsystem. A FlashCopy image is the logical equivalent of a complete physical copy, but you create it much more quickly and it requires less disk space.

Because a FlashCopy image is a host-addressable logical drive, you can perform backups by using the FlashCopy image while the base logical drive remains online and user-accessible. You can also write to the FlashCopy logical drive to perform application testing or scenario development and analysis. The maximum number of FlashCopy logical drives that are allowed is one-half of the total number of logical drives that are supported by your controller model.

VolumeCopy

The VolumeCopy feature is a firmware-based mechanism for replicating logical drive data within a storage array. This feature is designed as a system management tool for tasks such as relocating data to other drives for hardware upgrades or performance management, data backup, or restoring FlashCopy logical drive data.

You submit VolumeCopy requests by specifying two compatible drives. One drive is designated as the Source and the other as the Target. The VolumeCopy request is persistent so that any relevant result of the copy process can be communicated to you.

Note: The terms *VolumeCopy* and *Logical Drive Copy* are used interchangeably throughout this document, the Storage Manager Client program interface, and the online help.

Partition Expansion

You can use the Partition Expansion feature to present all of the defined logical drives in a storage subsystem to a SAN, through several different partitions. You can do this by *mapping* each logical drive to a LUN (0 - 31) in each of the defined partitions. A logical drive can be mapped to only one partition at a time. In a storage partition, the mapped logical drives are accessible only to the host server or servers whose host bus adapter host ports are defined in that storage partition. This feature enables the DS3000 storage subsystem to act as a multiple virtual storage subsystem to the host server.

Partition Expansion also enables multiple hosts that use different operating systems, with their own unique disk storage subsystem settings, to connect to the same DS3000 storage subsystem at the same time. This is known as a *heterogeneous host environment*.

Terms to know

For more information about terminology, see the online help in the Storage Manager Enterprise Management window or the Storage Subsystem Management window.

It is important to understand the distinction between the following two terms when you read this document.

Management station

A management station is a system that is used to manage a DS3000 storage subsystem. You can connect a management station to the storage subsystem by using either of the following methods:

- Through a TCP/IP Ethernet connection to the controllers in the storage subsystem (out-of-band method)
- Through a TCP/IP connection to the host-agent software that is installed on a host server, which in turn is directly attached to the storage subsystem through the input/output (I/O) path (in-band method)

For more information, see “Storage subsystem management methods” on page 6.

Host server

A host server is a server that is connected directly to a DS3000 storage subsystem through the I/O path. This server is used to perform the following tasks:

- Serve data (typically in the form of files) from the storage subsystem

- Function as a connection point to the storage subsystem for a remote-management station

Notes:

1. The terms *host* and *host server* are used interchangeably throughout this document.
2. A host server can also function as a management station.

Hardware requirements

Table 1 lists the hardware requirements for installing the Storage Manager 2 software.

Table 1. Hardware requirements for the Storage Manager 2 software

Hardware component	Requirements
Management station (for client software)	The management station computer requires the following hardware: • IBM System p™ or IBM BladeCenter® POWER™ server with 64-bit PowerPC® architecture • CD drive • Mouse or similar pointing device • System memory minimum of 256 MB (512 MB preferred) • Ethernet network interface adapter • Video adapter capable of a minimum of 800 x 600 resolution. • Monitor setting of 1024 x 768 pixels with 64 000 colors. The minimum display setting that is allowed is 800 x 600 pixels with 256 colors. **Important:** Some servers are not designed to run graphic-intensive software. If your server has difficulty displaying the Storage Manager user interface, you might have to upgrade the server video adapter or use the command-line interface (CLI).

For more information about host adapters, fabric switches, and managed hubs, see the following ServerProven® and interoperability Web sites:

http://www.ibm.com/servers/eserver/serverproven/compat/us/

http://www.ibm.com/systems/storage/disk/ds3000/ds3200/

http://www.ibm.com/systems/storage/disk/ds3000/ds3300/

http://www.ibm.com/systems/storage/disk/ds3000/ds3400/

AIX software requirements

The Storage Manager software requirements and the IBM AIX operating system requirements are described in the following sections.

Storage Manager software requirements

Table 2 on page 5 lists the disk space and administrator privileges that are required to install the Storage Manager software in an AIX environment.

Note: The disk space that is listed in Table 2 is for the packages after the installation. During the installation of each software package, you might need up to 170 MB of free disk space.

Table 2. Installation requirements for AIX by software package

Package	Disk space requirement for AIX
SMclient	130 MB
MPIO	2 MB
SMagent	1 MB 50 MB (if SMclient is not installed)
SMutil	2 MB

AIX operating system requirements

The management station or host server requires one of the following operating systems.

Important: For the latest supported operating system and host software versions, see the Storage Manager readme files and see http://www.ibm.com/servers/eserver/serverproven/compat/us/.

- IBM AIX version 5.2
- IBM AIX version 5.3
- IBM AIX version 6.1

Linux on POWER software requirements

The Storage Manager software requirements and the Linux operating system requirements are described in the following sections.

Storage Manager software requirements

Table 3 lists the disk space that is required for installing the Storage Manager software in a Linux environment.

Table 3. Installation requirements for Linux by software package

Package	Disk space requirement
Storage Manager Client (SMclient)	165 MB in /opt, 2 MB in /tmp, 1 MB in /var
Linux MPP	2 MB
Storage Manager Agent (SMagent)	2 MB in /opt
Storage Manager Utility (SMutil)	7 MB in /opt

Linux operating system requirements

The following kernel levels are required for Linux operating systems:

- 2.6 kernel (RHEL 4, SLES 9, and SLES 10) for RDAC Revision B
- 2.6 kernel (RHEL 5 and SLES 10 SP1) for RDAC Revision C

For the latest supported Linux operating systems, details about device driver levels, and instructions for multipath configurations for Linux support, see the Storage Manager software readme files.

Storage subsystem management methods

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

- The host-agent (in-band) management method (DS3300 and DS3400 only)
- The direct (out-of-band or Ethernet) management method

If you establish management connections from multiple management stations or hosts to the same DS3000 storage subsystem, it is best practice to have only one active management connection to the DS3000 storage subsystem at a time, even though eight simultaneous active management sessions are supported in a DS3000 storage subsystem.

Host-agent (in-band) management method

When you are using the host-agent (in-band) management method, you manage the storage subsystems through the I/O path to the host. The management information can either be processed in the host or be passed to the management station through the network connection, as shown in Figure 1.

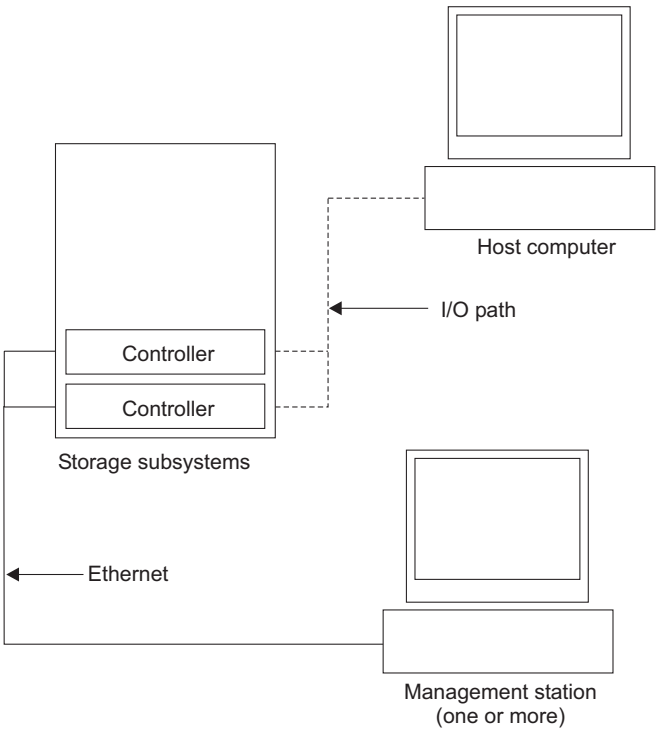

Figure 1. Host-agent (in-band) managed storage subsystems

Managing storage subsystems through the host-agent has the following advantages:

- You do not have to connect Ethernet cables to the controllers.
- You do not need a Dynamic Host Configuration Protocol/Bootstrap Protocol (DHCP/BOOTP) server to connect the storage subsystems to the network.
- You do not have to perform the controller network configuration tasks that are described in Chapter 2, “Preparing for installation,” on page 15.

- You must specify only a host name or IP address for the host instead of for the individual controllers in a storage subsystem. Storage subsystems that are attached to the host are automatically discovered, if the host-agent service is installed and running.

Managing storage subsystems through the host-agent has the following disadvantages:

- The host-agent requires a special logical drive, called an *access logical drive,* to communicate with the controllers in the storage subsystem. Therefore, you are limited to configuring one fewer logical drive than the maximum number that is allowed by the operating system and the host adapter that you are using.
- If the connection through the I/O path is lost between the host and the storage subsystem, you cannot manage or monitor the storage subsystem.

Note: The access logical drive is also referred to as the *Universal Xport Device*.

Important: If the maximum number of logical drives are configured in the host, either use the direct-management method or give up a logical drive for use as the access logical drive. For more information, see “Number of supported logical drives” on page 14.

Required host-agent (in-band) management activities

To establish host-agent (in-band) management connections, complete the following general tasks:

1. Make the I/O connections between the DS3000 storage subsystem and the host server.
2. Install TCP/IP on the host server and establish a host IP address. For instructions for host-managed systems, see “Overview of network installation tasks” on page 16.
3. Install the host bus adapter device driver, SMclient, RDAC (MPIO for AIX, Linux MPP for Linux), and SMagent on the host server.
4. Make sure that the Storage Manager Agent Service starts successfully:

   For AIX and Linux: Type `ps -elf | grep SMagent`
5. Open the Storage Manager Enterprise Management window and add the IP address of the host server to the management domain. All DS3000 devices that are recognized by the host server are displayed in the Enterprise Management window.

Direct-management (out-of-band) method

When you use the direct-management (out-of-band) method, you manage storage subsystems directly over the network through a TCP/IP 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. Each of the storage subsystem controllers must be connected, through a cable that is connected to the RJ-45 connector, to an Ethernet network, as shown in Figure 2 on page 8.

The following illustration shows a system in which storage subsystems are managed by using the direct-management (out-of-band) method.

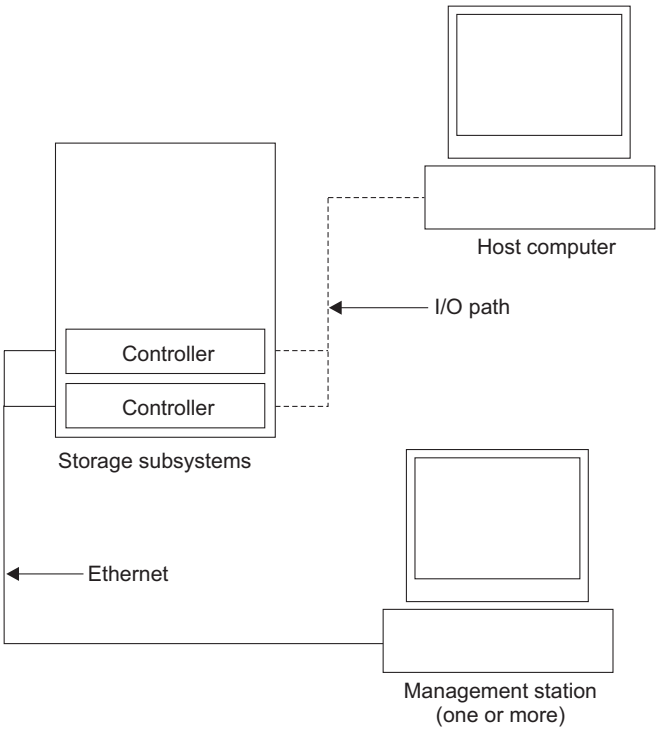

Figure 2. Direct-management (out-of-band) storage subsystems

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

- The Ethernet connections to the controllers enable you to manage storage subsystems that are connected to a host running any of the operating systems that Storage Manager supports through an SMclient management station.
- You do not need an access logical drive to communicate with the controllers. You can configure the maximum number of logical drives that are supported by the operating system and the host bus adapter that you are using.
- You can manage and troubleshoot the storage subsystem when there are problems with the Fibre Channel, Internet SCSI (iSCSI), or SAS links.

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

- In a dual-controller storage subsystem configuration, you need two Ethernet cables to connect the storage subsystem controllers to a network.
- When you add devices, you must specify an IP address or host name for each controller.
- Network preparation tasks are required. For a summary of the preparation tasks, see Chapter 2, “Preparing for installation,” on page 15.

Important:

- To avoid DHCP/BOOTP server and network tasks, use the controller default TCP/IP addresses or assign static IP addresses to the controller (see Table 4 on page 9).
- Create a separate private network aside from the main production network for managing your DS3000 storage subsystems.

Defining controller TCP/IP addresses

The following table shows the default settings for the storage subsystem controllers.

Table 4. Default controller TCP/IP addresses

Controller	IP address	Subnet mask
A	192.168.128.101	255.255.255.0
B	192.168.128.102	255.255.255.0

To change the DS3000 storage subsystem controller default TCP/IP address, you must first make a direct-management connection to the storage subsystem by using the default TCP/IP addresses.

To change the controller default TCP/IP address by using the Storage Manager Client program, complete the following steps:

1. In the Enterprise Management window, double-click the storage subsystem for which you want to define new controller TCP/IP addresses. The Subsystem Management window opens.
2. Click the **Tools** tab.
3. On the Tools page, click **Change Network Configuration**. The Change Network Configuration window opens.

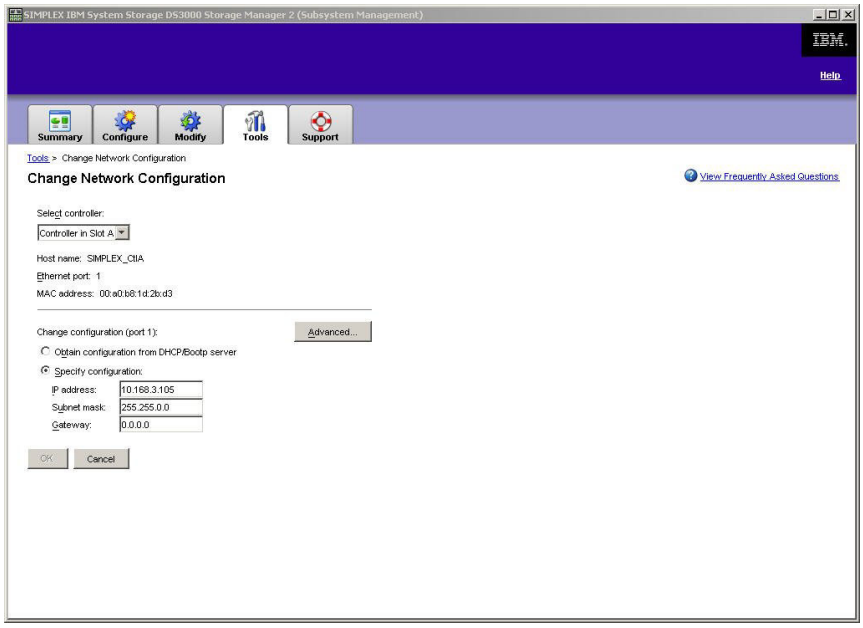

Figure 3. Change Network Configuration window

4. Select the applicable controller tab and enter the applicable TCP/IP address information for both controllers A and B.
5. Click **OK** to apply the new TCP/IP addresses.

   Important: Make sure that the new TCP/IP address information is correct because you will not be able to make a direct management connection to the storage subsystem again by using the controller default TCP/IP addresses after they are changed to different values.
6. Close the Subsystem Management window.

7. In the Enterprise Management window, delete the storage subsystem entry that contains the default TCP/IP addresses.
8. Wait approximately 5 minutes, and then re-establish the management connection to the storage subsystem by using the new TCP/IP addresses.

   Note: You might want to ping the new addresses first to make sure that there is a good network connection to the controllers before you add the storage subsystem into the Enterprise Management window.

Required direct-management (out-of-band) management activities

To establish a direct-management connection, complete the following steps:

1. Make the Ethernet connections between the DS3000 storage subsystems and the Ethernet network.
2. Follow the instructions in “Overview of network installation tasks” on page 16 for direct-managed systems.
3. Install SMclient.
4. Start the Storage Manager software and add the IP addresses for the DS3000 controllers (A and B). If you have multiple DS3000 storage subsystems, add each IP address to the domain.

Reviewing a sample network

The following illustration shows an example of a network that contains both a direct-managed storage subsystem (Network A) and a host-agent-managed storage subsystem (Network B).

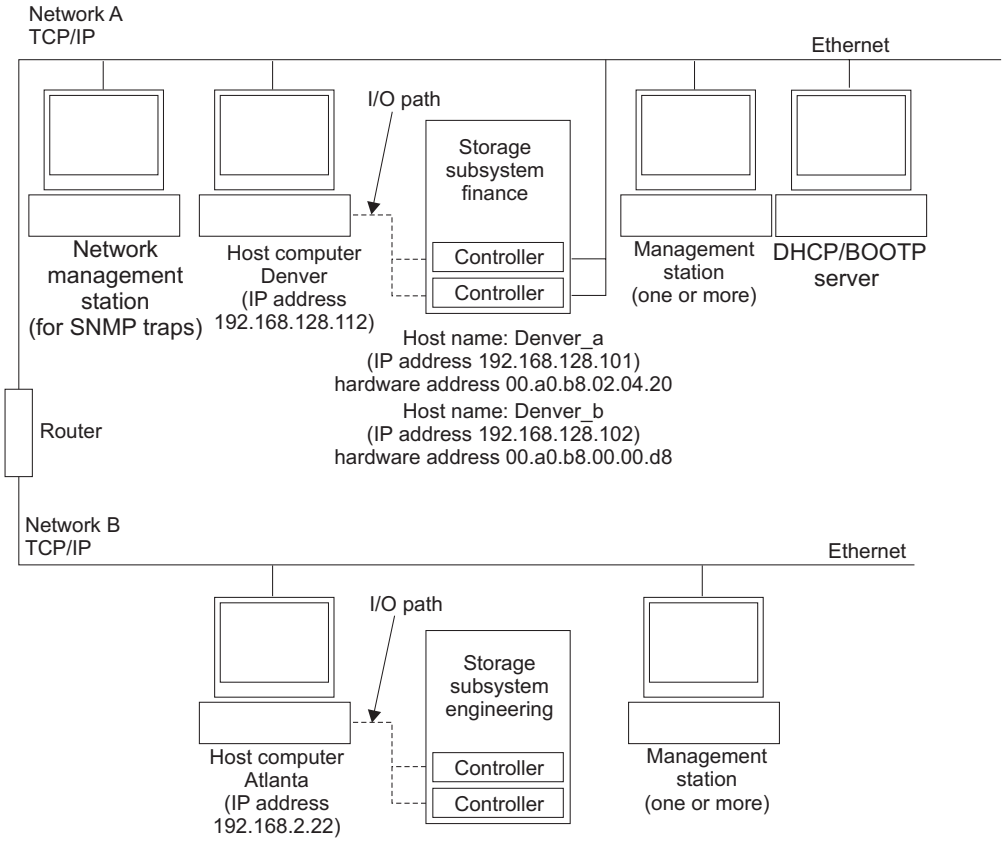

Figure 4. Sample network using direct-managed and host-agent managed storage subsystems

Direct-managed storage subsystem

In Figure 4, Network A is a direct-managed storage subsystem. Network A contains the following components:

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

Note: If you use the controller static TCP/IP addresses or default TCP/IP addresses, you do not have to setup the DHCP/BOOTP server.

Host-agent-managed storage subsystem

In Figure 4 on page 11, Network B is a host-agent-managed storage subsystem. Network B contains the following components:

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

Creating a SAN-attached configuration (Fibre Channel)

This section contains the following topics:

- Steps for setting up a SAN-attached configuration
- Connecting host bus adapters in a Fibre Channel switch environment

Setting up a SAN-attached configuration

To set up a SAN-attached configuration, complete the following steps:

1. Connect the host bus adapters to the switch or switches.
2. Connect the DS3000 storage subsystems to the switch or switches.
3. Set the required zones on the Fibre Channel switch or switches, if applicable.

   Note: For information about zoning and enabling zones, see the documentation that comes with the switch.
4. Configure the system, and then verify the configuration by logging in to the switch to view the system.

For more information about setting up a DS3000 storage subsystem, see the *Quick Installation Guide* and the other documentation on the support CD that comes with the storage subsystem.

Standard (noncluster) configuration

You can install the Storage Manager software in a standard (noncluster) configuration.

Figure 5 shows a sample standard (noncluster) configuration.

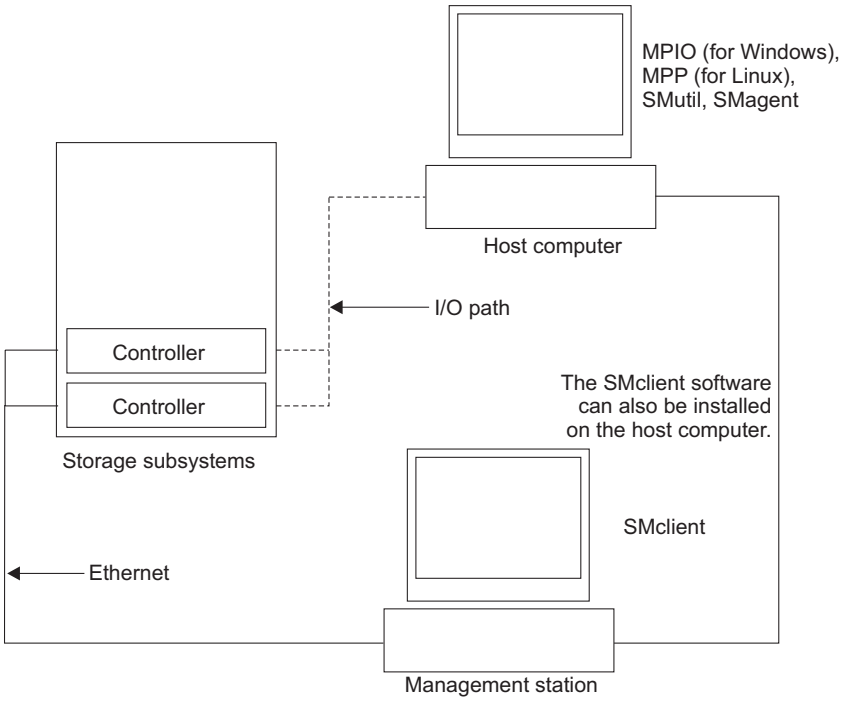

Figure 5. Sample standard (noncluster) configuration

Number of supported logical drives

The supported logical drive limits for AIX and Linux on POWER are shown in the following table.

Table 5. Maximum number of logical drives per storage subsystem

Maximum logical drives	Maximum logical drives per storage unit partition	Maximum FlashCopy logical drives	Maximum Logical Drive Copy logical drives
256	32 (See note 1.)	128	255

Notes:

1. A maximum of four partitions is supported for base DS3000 models. To upgrade from 4 to 16 partitions, you must purchase the DS3000 Partition Expansion License.
2. Each standard logical drive supports a maximum of four FlashCopy logical drives.
3. The host-agent management method uses a special logical drive, called an *access logical drive*, to communicate with the controllers on the storage subsystem. The access logical drive uses one of the available LUNs. Therefore, managing storage subsystems with the host-agent software limits you to one fewer LUN than the maximum number that is supported by the operating system and the host adapter. The access logical drive is also referred to as the *Universal Xport Device*.

Chapter 2. Preparing for installation

This chapter provides information to help you plan the installation of the Storage Manager software.

See the applicable sections in this document to install the Storage Manager 2 Client and host software for a particular operating system. When you have completed the Storage Manager 2 software component installation procedures, perform the procedures in Chapter 5, “Completing the Storage Manager software installation and configuration,” on page 75 and Chapter 6, “Completing storage-management tasks,” on page 107.

Planning and preparing for an AIX installation

Table 6 shows the installation sequence for the Storage Manager software components and provides information to help you plan for the installation of the Storage Manager software for AIX.

Table 6. Software component installation sequence in an AIX environment

System type	Software installation sequence
Management station (and any host server in use as a management station)	Storage Manager Client (SMclient)
Host server	1. IBM AIX 5.2, AIX 5.3, or AIX 6.1 and the required APARS 2. IBM host bus adapter device driver 3. IBM AIX MPIO PCM (required) Multipath I/O (MPIO) path-control module (PCM) is installed as the failover driver for AIX. 4. Storage Manager host software package, which can install the following components: • SMclient (optional) SMclient is required if you manage the storage subsystem from the host server. • MPIO (required) MPIO is installed as the RDAC failover driver for AIX. • SMagent (optional) SMagent is required if you use the host-agent (in-band) management method. The direct (out-of-band) management method is preferred in an AIX environment. • SMutil (optional) SMutil contains the hot_add, SMdevices, and SMrepassist utilities.

Planning and preparing for a Linux installation

Table 7 shows the installation sequence for the Storage Manager software components and provides information to help you plan for the installation of the storage management software for Linux.

Table 7. Software component installation sequence in a Linux environment

System type	Software installation sequence
Management station (and any host server in use as a management station)	Storage Manager Client (SMclient)
Host server and cluster node	1. IBM host bus adapter device driver (Linux version) 2. Linux MPP (required) Linux MPP is installed as the RDAC failover driver for Linux. 3. Storage Manager host software package, which can install the following components: • SMclient (optional) SMclient is required if you manage the storage subsystem from the host server. • SMagent (optional) SMagent is required if you use the host-agent (in-band) management method. You must first install Linux MPP (see step 2). The direct (out-of-band) management method is preferred in a Linux environment. • SMutil (required) SMutil contains the hot_add, SMdevices, and SMrepassist utilities. 4. Host bus adapter management software (optional)

Note: Linux MPP is provided separately from the Storage Manager host software package. There are two Linux MPP packages, one for each type of kernel:

- 2.6 kernel (RHEL 4, SLES 9, and SLES 10) for RDAC Revision B
- 2.6 kernel (RHEL 5 and SLES 10 SP1) for RDAC Revision C

Overview of network installation tasks

To begin installing the storage management software, make sure that the network components are set up and operating correctly, and that you have the host and controller information that is required for the software to operate correctly.

The following sections provide an overview of the tasks that are involved in setting up a network installation for direct-managed and host-managed storage subsystems.

Direct-managed storage subsystems: Network installation overview

Setting up a network installation of a direct-managed system involves the following general tasks:

1. Install all hardware components (host servers, storage subsystems, and cables) that you want to connect to the network. For installation information, see the documentation that comes with the specific hardware device.

2. Establish a naming convention for the storage subsystems that are connected to the network. For more information, see “Step 1: Naming the storage subsystems” on page 18.
3. Record the storage subsystem names and management types in Appendix A, “Storage subsystem and controller information record,” on page 113. An example of a completed information record is shown in Table 8 on page 18.

   Note: Throughout the remaining steps, return to Appendix A, “Storage subsystem and controller information record,” on page 113 to record information such as hardware Ethernet addresses and IP addresses.
4. If you are using a default controller IP address, go to step 6. Otherwise, obtain the TCP/IP address and host name for each of the controllers in the storage subsystems on the network from the network administrator. See “Step 3: Assigning IP addresses to hosts and controllers” on page 19.
5. Use controller static IP addresses (preferred); otherwise, set up a DHCP/BOOTP server to provide network configuration information for a specific controller.
6. Make sure that the TCP/IP software is installed. For more information, see “Step 4: Verifying the TCP/IP protocol and setting up the host or DNS table” on page 19.
7. Power-on the devices that are connected to the network.

Host-managed systems: Network installation overview

The following steps provide an overview of the tasks that are involved in setting up a network installation of a host-managed system:

1. Install all hardware components (host servers, storage subsystems, and cables) that you want to connect to the network. For installation information, see the documentation that comes with the specific hardware device.
2. Establish a naming convention for the storage subsystems that are connected to the network. For more information, see “Step 1: Naming the storage subsystems” on page 18.
3. Record the storage subsystem names and management types in Appendix A, “Storage subsystem and controller information record,” on page 113. An example of a completed information record is shown in Table 8 on page 18.

   Note: Throughout the remaining steps, return to Appendix A, “Storage subsystem and controller information record,” on page 113 to record information such as hardware Ethernet addresses and IP addresses.
4. Obtain the IP address and host name of the host server on which the host-agent software will run from the network administrator. See “Step 3: Assigning IP addresses to hosts and controllers” on page 19.

   Note: Host-managed storage subsystems are supported with the DS3300 and DS3400 only.
5. Make sure that the TCP/IP software is installed. For more information, see “Step 4: Verifying the TCP/IP protocol and setting up the host or DNS table” on page 19.
6. Power-on the devices that are connected to the network.

Step 1: Naming the storage subsystems

As you set up your network, decide on a naming convention for the storage subsystems. When you install the storage management software and start it for the first time, all storage subsystems in the management domain are displayed as `<unnamed>`. Use the Subsystem Management window to rename the individual storage subsystems.

Consider the following factors when you name storage subsystems:

- There is a 30-character limit. All leading and trailing spaces are deleted from the name.
- Use a unique, meaningful naming scheme that is easy to understand and remember.
- Avoid arbitrary names or names that might quickly lose their meaning.
- The software adds the prefix `Storage Subsystem` when storage-subsystem names are displayed. For example, if you name a storage subsystem Engineering, it is displayed as `Storage Subsystem Engineering`.

Step 2: Recording the storage subsystem names

After you decide on a naming scheme, record the storage subsystem names in the blank information record form in Appendix A, “Storage subsystem and controller information record,” on page 113).

Table 8 shows a sample information record for the network shown in Figure 4 on page 11. This network contains storage subsystems that are managed by using both the direct-management and host-agent-management methods.

Table 8. Sample information record

Storage subsystem name	**Management method**	**Controllers: IP addresses and host name**		**Host-IP address and host name**
		Controller A	**Controller B**	
Finance	Direct	IP address = 192.168.128.101	IP address = 192.168.128.102	
		Host = Denver_a	Host = Denver_b	
Engineering	Host-agent			IP address = 192.168.2.22
				Host = Atlanta

Step 3: Assigning IP addresses to hosts and controllers

If you plan to manage your storage subsystem by using the out-of-band management (direct-management) method, assign a unique IP address for every controller in all storage subsystems on the network. Record the IP address for each controller in the information record (Appendix A, “Storage subsystem and controller information record,” on page 113).

Note: To avoid performing DHCP/BOOTP server and network tasks, use the controller default TCP/IP addresses or assign static IP addresses to the controllers. If you are not able to use the controller default TCP/IP address to assign the static TCP/IP address, see the information in “Direct-management (out-of-band) method” on page 7.

You can establish the out-of-band management connection to the DS3000 storage subsystem by using the default IP addresses of 192.168.128.101 for controller A and 192.168.128.102 for controller B. You can change the IP addresses by using the Storage Manager Client Subsystem Management window.

For host-agent management only: If you plan to manage the DS3300 or DS3400 storage subsystem by using the host-agent management method, assign an IP address for each host on which you will install the host-agent software. Record the IP address for each host in the information record in Appendix A, “Storage subsystem and controller information record,” on page 113; then, go to “Step 4: Verifying the TCP/IP protocol and setting up the host or DNS table.”

Step 4: Verifying the TCP/IP protocol and setting up the host or DNS table

The default installation process installs the TCP/IP software package. To make sure that TCP/IP software package was installed successfully, open a shell prompt and type `ifconfig`. If you see an eth0 or eth1 interface, the IP protocol is enabled.

To set up the host or DNS table, complete the following steps. Make sure that the host names for the controller match the IP addresses for the controllers.

1. Update either the host table or the DNS table to specify a host name to associate with an IP address. If you do not have a DNS table, edit the host table that is in the /etc/hosts file. Your directory might be different if the operating system is not installed on the root.
2. To manage storage subsystems through a firewall, configure the firewall to open port 2463 to TCP data.
3. Attempt to ping one of the controllers that are listed in the /etc/hosts file.

The following example shows a Linux host file.

```
# Do not remove the following line, or various programs
# that require network functionality will fail.

127.0.0.1 rh71_8500 localhost.localdomain
102.54.94.97 rhino.acme.com x.acme.com
```

Chapter 3. Installing and configuring the Storage Manager software on AIX host systems

This chapter describes how to install the Storage Manager software in an AIX operating-system environment.

Hardware and software requirements for the AIX host system

This section provides information about the hardware and software requirements for the AIX host system.

Note: For the latest installation requirements and user information about the Storage Manager software, AIX file sets, and hardware components, see the Storage Manager readme file for AIX that is on the *System Storage DS3000 Support Software for AIX* CD (AIX host kit support CD).

Hardware requirements

The following minimum hardware is required.

Note: For the most current host support information, see the latest DS3000 interoperability matrix.

DS3200 support:

- IBM BladeCenter H chassis or IBM BladeCenter E chassis with two IBM BladeCenter SAS Connectivity Modules
- IBM BladeCenter JS12, IBM BladeCenter JS21, and IBM BladeCenter JS22, using the IBM BladeCenter SAS Expansion Card (CFFv)
- One BladeCenter unit can connect to only one DS3200 storage subsystem.

  Note: To configure a BladeCenter unit and a DS3200 storage subsystem, connect host port 1 on a DS3200 controller to a BladeCenter SAS Connectivity Module in a BladeCenter unit. Connect the other host port 1 on the second DS3200 controller to a second BladeCenter SAS Connectivity Module in a BladeCenter unit. No other external devices can be connected to the SAS connectivity modules.

DS3400 support:

- Any IBM System p or POWER server using the IBM POWER4™, POWER5™, or POWER6™ 64-bit PowerPC architecture that supports one of the following Fibre Channel host bus adapters:
  - FC 6239
  - FC 5716 / 1957 / 1977
  - FC 5758 / 1905
  - FC 5759 / 1910
  - FC 5773
  - FC 5774
- IBM BladeCenter JS21 and IBM BladeCenter JS22, using the IBM @server BladeCenter Fibre Channel expansion cards

Software requirements

The following operating system software is supported:

- AIX 5.2

  The operating system must be at version 5200-10-04-0750 or later.

  Note: The DS3200 storage subsystem is not supported on AIX 5.2.
- AIX 5.3

  The operating system must be at version 5300-08-03-0831 or later.
- AIX 6.1

  The operating system must be at version 6100-01-02-0834 or later.

Note: The required minimum APARS are listed in the Storage Manager software readme file for AIX.

AIX restrictions

The following restrictions apply to AIX configurations:

SAN and connectivity restrictions

- AIX has the following limitations on disk volume sizes:
  - 1 TB, on 32-bit AIX version 5.2
  - 2 ZB, on 64-bit AIX version 5.2 (2 TB when using LVM bad block relocation)

  Where larger volumes are supported, the AIX boot logical volume must reside within the first 2 TB of the volume. Hardware devices or vendor software might further limit volume sizes. For more information, contact your hardware or software vendor.
- Other storage devices, such as tape devices or other disk storage, must be connected through separate HBAs and SAN zones.

Restrictions when starting the system

- When you start the system from a DS3000 device, both paths to the DS3000 storage subsystem must be up and running.
- The system cannot use path failover during the AIX startup process. After the AIX host is started, failover operates normally.

For the best performance if you are starting the system from a SAN, do not use a logical drive created on a SATA array.

Partitioning restrictions

- The maximum number of partitions per AIX host, per DS3000 storage subsystem, is three.
- All logical drives that are configured for AIX must be mapped to the AIX host group (not AIX with Veritas DMP).

  For more information, see “Configuring host access” on page 83.
- 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.
- AIX can support only a single DS3200 partition.
- AIX can support only dual-controller DS3200 models and there must be an I/O path from the host to each DS3200 controller.

Installing the DS3000 Storage Manager software

This section describes the procedures for installing the Storage Manager software.

Preparing for software installation

Before you install the Storage Manager software, make sure that you have the following items:

- For the AIX operating system, make sure that you are running the operating system-required maintenance level. (See the readme file for the lists of required maintenance levels and additional fixes.)
- HBA device drivers
- The most up-to-date controller firmware for your particular storage subsystem, but at a minimum the following firmware and NVSRAM versions are required:

  DS3200:

Controller firmware:	07.35.41.00
Controller NVSRAM:	N1726D32LR335V02 (single controller) N1726D320R335V06 (dual controller)

  DS3400:

Controller firmware:	07.35.41.00
Controller NVSRAM:	N1726D34LR335V02 (single controller) N1726D340R335V05 (dual controller)

- IP addresses for the RAID controllers (for out-of-band management only)
- Additional documentation for switches and HBAs, if needed
- Applicable host software kit. The host software kit grants you permission to attach host servers using the applicable operating system for the DS3000 storage subsystem. The kit includes the most up-to-date DS3000 software and controller firmware for your particular storage subsystem.
- AIX host kit support CD, which includes the most up-to-date DS3000 software and controller firmware for your particular storage subsystem

  For the latest controller firmware, see http://www.ibm.com/servers/eserver/support/fixes/.

Failover drivers

An AIX host system requires the AIX multipath I/O (MPIO) failover driver for path redundancy. The failover driver monitors I/O paths. If a component failure occurs in one of the I/O paths, the failover driver reroutes all I/O to another path.

Note: The AIX MPIO driver files are not included on the AIX host kit support CD. To install them, you must download the files from http://www.ibm.com/systems/support/.

Make sure that the AIX device driver that supports the DS3000 is installed. The Multi Path I/O driver (MPIO) is the driver that is required to run the DS3000 on AIX systems. Verify that the driver is installed by typing the following command:

```
lslpp -i devices.common.IBM.mpio.rte
```

If the device driver is not installed, the command returns ″Fileset devices.common.IBM.mpio not installed.″ Obtain the required AIX installation media and use the AIX SMIT utility to install the MPIO package. After the installation of the MPIO package, make sure that the operating system version still reports the required operating system version using the **oslevel -s** command.

Preparing for a SAN attachment (DS3400 only)

AIX supports either a direct-attached or a SAN-attached configuration to the DS3400 storage subsystem.

The AIX failover driver that supports the DS3400 is the Multi Path I/O driver (MPIO). The MPIO driver provides flexibility in the connectivity methods that can be used when connecting supported devices.

Single host bus adapter configurations are allowed. If you have a single host bus adapter in a SAN-attached configuration, the controller must have a Fibre Channel connection to the host bus adapter through a switch, and both controllers must be within the same SAN zone as the host bus adapter for dual-controller storage configurations.

The following examples are a small subset of the connectivity configurations that are supported using the DS3400. For more configuration examples, see the section "Connecting hosts to the DS3400" in the *IBM System Storage DS3400 Storage Subsystem Installation, User's, and Maintenance Guide*.

Attention: If a path fails in a single host bus adapter configuration, a loss of data access can occur.

The best practice is to use 2 to 4 adapters providing 4 to 8 paths.

Example 1

Create one zone that contains two Fibre Channel HBA ports, one Fibre Channel port from controller A and one Fibre Channel port from controller B of the DS3400.

This configuration provides four paths to hdisks. Two of the paths that are associated with the preferred storage controller will service I/Os when the DS3400 is Optimal. The other two paths are used if the preferred controller is not accessible by the host.

Note: The hdisk attribute "algorithm" should be set to round_robin. If the attribute is set to round_robin, the hdisk attribute "reserve_policy" must be set to no_reserve, pr_exclusive, or pr_shared.

Example 2

Create separate zones for the connection between the HBA port and one port of the DS3400 controller. Two zones are required for a dual-controller storage configuration. One zone contains an HBA port and controller port from controller A. The other zone contains a different HBA port and controller port from controller B.

This configuration provides two paths to hdisks.

Note: Attributes "algorithm" and "algorithm of hdisks" do not need to be altered from default values.

Example 3

Create one zone that contains two Fibre Channel HBA ports and all four Fibre Channel ports from the DS3400.

This configuration provides eight paths to hdisks. Four of the paths that are associated with the preferred storage controller will service I/Os when the DS3400 is Optimal. The other four paths are used if the preferred controller is inoperable.

Note: The hdisk attribute ″algorithm″ should be set to round_robin. When this attribute is set to round_robin the hdisk attribute ″reserve_policy″ must be set to no_reserve, pr_exclusive, or pr_shared.

Preparing for attachment to a BladeCenter unit (DS3200 only)

AIX supports the attachment of only one dual-controller DS3200 to a BladeCenter unit.

The attachment requires two SAS connectivity modules. One SAS connectivity module is connected by a single cable to host port 1 of controller A. The second SAS connectivity module is connected by a single cable to host port 1 of controller B. No other external devices can be connected to the SAS connectivity modules or to host ports 2 or 3 of the DS3200 controller modules.

Installing the Storage Manager software using the installation wizard

The DS3000 Storage Manager software installation wizard installs the following software packages on the AIX host:

- SMruntime
- SMclient
- SMagent
- SMutil

Requirements

- DS3000 Storage Manager version 02.70, or later
- xservices must be operational

If you are installing Storage Manager 02.70, 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 the Storage Manager software in console mode” on page 27.

Note: If you use the wizard to install the Storage Manager software, you might need to increase the size of some partitions to accommodate the installation.

To install the DS3000 Storage Manager software using the installation wizard, complete the following steps. Adjust the steps as necessary for your specific installation.

1. Copy the following file from the AIX host kit support CD or download the file from the DS3000 support Web site, to a directory on your system:

   SMIA-AIX-02.70.xx.xx.bin
2. Open the file on your system. The Storage Manager Installation wizard 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

Note: During the installation, you will see the question **Automatically Start Monitor?** This refers to the Event Monitor service. If you want to enable automatic ESM firmware synchronization, the Event Monitor must be enabled. To enable the Event Monitor, select **Automatically Start Monitor**.

The DS3000 Storage Manager software is installed on the system.

Installing the Storage Manager software in console mode

To install the Storage Manager software without the graphical interface, complete the following steps:

1. From a command prompt, run the Storage Manager installation package while using the - i console parameter to force the package to install in a non-graphical mode. For Example, in Linux, the command is:

   ```
   # sh SMIA-LINUXPPC-02.70.A5.15.bin – i console
   ```

2. Select the location (the default value is English) and press Enter.
3. Read the End User License Agreement (EULA). To continue, press **Y** (for Yes) to accept the EULA.
4. Choose the installation type, and then press Enter:
   - Typical
   - Management Station
   - Host
   - Customize
5. Review the Pre-Installation Summary, and then press Enter. The installation is completed.
6. Press Enter to exit the installer.

Configuring storage subsystems

If you have purchased a Storage Partitioning premium feature, make sure that the premium feature is enabled. For more information, see Chapter 7, “Enabling and using premium features,” on page 109.

To configure the Storage Manager software for an AIX system, complete the procedures in the following sections. You can configure the Storage Manager software from the Storage Manager client software running on either an AIX or a non-AIX system.

Adding storage subsystems to the Storage Manager Client

To add storage subsystems to the Storage Manager Client (SMclient), complete the following steps:

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

   Note: For information about assigning IP addresses to the controllers, see “Step 3: Assigning IP addresses to hosts and controllers” on page 19.

2. After the storage subsystem is configured on the network, start SMclient on the host server by selecting the Storage Manager icon or 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 is displayed in the Enterprise Management window.

      Note: Failure to add both controllers results in a partially-managed system, where all functions might not be available.
4. Double-click the name of the storage subsystem to open its Subsystem Management window.

Upgrading controller firmware and NVSRAM

For detailed procedures for downloading the firmware updates, see “Downloading controller, NVSRAM, ESM, and hard disk drive firmware” on page 94.

Creating storage partitions

Before you create storage partitions, make sure that you have created a physical connection between the host and the storage subsystem controllers, and that you have connected and zoned the SAN switch (if any). If you have not completed these connections, the Storage Manager software is *not* able to list the WWPNs of the HBAs during the procedures in the following sections.

Recording the WWPN of the HBAs on a server

Type the following commands for each HBA port that you want to use for the installation. This documents the WWPN that is displayed in the Storage Manager Client during the host definition process.

```
lscfg -vpl fcsX|grep Network

# lscfg -vpl fcs0|grep Network

Network Address.............10000000C94BC2A3
```

Record this information to use during the procedure for defining the host and host port.

Note: As of the date of this document, AIX does not enable the SAS adapter WWPN to be displayed. The UUID that is displayed in the BladeCenter unit hardware VPD menu for the SAS expansion card can be used to determine the host port identifier for the first physical connection. To determine the second host port identifier, the first host port identifier is incremented by 1. For example, if the first host port identifier ends in 5432, the second host port identifier is updated so that it ends in 5433.

Running the cfgmgr command

Run the AIX command **cfgmgr -v**. This causes the HBA ports to log into the storage ports. This enables the Storage Manager software to display the WWPN during the host definition process.

Defining the host and host port

To define the host and host ports by using the Configure Host Access wizard, complete the following steps:

1. Select the **Configure** tab on the Subsystem Management window and select one of the following methods:
   - Create Host Access (Automatic)
   - Create Host Access (Manual)
2. Select **Create Host Access (Automatic)** if you have added a new host that is also running the host agent on the system. Otherwise go to step 5.
3. If any host are automatically configured, they are displayed in the left column as an Available Host. Select a host from the list of available hosts and click **Add**.
4. After all the available hosts are moved to the **Select Host** column, click **OK**. The host and host ports are defined and you have completed this procedure.
5. Select **Create Host Access (Manual)** to manually define the host and host port information.
6. From the Specific Host Name and Host Type window, enter the host name and then select **AIX** as the host type from the host type list. Click **Next**.
7. From the Specify HBA Host Port window, in the left pane, select the correct WWPN of the HBA host port that is used to define the host. Click **Add**.
8. After all the host ports are added to the **Selected HBA host port** column, click **Next**.
9. From the Specify Host Group window, select one of the following options:
   - Select **No** if the host will *not* share access to the same logical drive with other hosts; then, click **Next**.
   - Select **Yes** if the host will share access to the same logical drive with other hosts. Either enter a New Host Group name or select an existing host group from the list; then, click **Next**.
10. Review the host definition details, and then click **Finish**.

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.

To define a host group, complete the following steps:

1. Click the **Configure** tab on the Subsystem Management window.
2. Select **Create Host Group**.
3. Enter the new host group name (for example, AIX) and select the hosts from the left column. Click **Add** to add them to the host group.
4. When all the hosts for the host group are added, click **OK**.
5. A message is displayed indicating that the new host group is created. Click **OK**.

Mapping LUNs to a storage partition

This section contains the procedures for mapping and adding LUNs to a host or host group. These steps can be used to map LUNs to a new partition or add LUNs to an existing partition.

To map LUNs, complete the following steps:

1. Select the **Configure** tab on the Subsystem Management window.
2. Select **Create Host-to-Logical Drive Mapping**.
3. From the Select Host window, select the host or host group that the logical drives will be mapped and click **Next**.
4. From the Select Logical Drives window, select the logical drives that will be mapped to the host or host group that was selected in step 3.

   Note: If all the logical drives will be mapped to the same host or host group, then click the **Select all logical drives** check box.

   After the logical drives are selected, assign a logical unit number (LUN) from 0 to 31, to the logical drives and click **Finish**.
5. The Host-to-Logical Drive Mapping progress window is displayed. When the mappings are completed, click **OK**.
6. The mappings are completed. Select **Yes** to map additional logical drives or **No** to complete the mappings wizard.

Identifying devices on AIX hosts

The MPIO driver creates the following devices that represent the DS3000 storage subsystem configuration:

dac The disk array controller (dac) devices represent a controller within the storage subsystem. Dac devices are only needed for in-band management. If access LUNs are not mapped to the system partition through the Storage Manager software, no dac devices are discovered.

hdisk Each hdisk device represents an individual logical drive on the AIX system. Logical drives that are owned by the DS3000 storage subsystems can be identified using the following command:

```
lsdev -Cc disk|grep DS3K
```

Performing initial device discovery

Before you begin: Make sure that the DS3000 storage subsystem is set up correctly, any applicable Fibre Channel switch zoning is complete, LUNs are assigned to the host, and the AIX software and firmware requirements are met.

To perform the initial device discovery, complete the following steps:

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 the **cfgmgr** command.
2. Type the following command:

   ```
   # lsdev -Cc disk | grep DS3K
   ```
3. Examine the output of the **lsdev -Cc disk | grep DS3K** command to be sure that the AIX MPIO software recognizes the correct number of DS3000 logical drives.

   The following example shows the output of the **lsdev -Cc disk | grep DS3K** command for a set of DS3400 LUNs:

```
# lsdev -Cc disk|grep DS3K
hdisk3   Available 00-08-01     MPIO Other DS3K Array Disk
hdisk4   Available 00-08-01     MPIO Other DS3K Array Disk
hdisk5   Available 00-08-01     MPIO Other DS3K Array Disk
hdisk6   Available 00-08-01     MPIO Other DS3K Array Disk
hdisk7   Available 00-08-01     MPIO Other DS3K Array Disk
hdisk8   Available 00-08-01     MPIO Other DS3K Array Disk
hdisk9   Available 00-08-01     MPIO Other DS3K Array Disk
```

Verifying the installation and configuration

After you have performed the initial device identification, complete the following procedures to verify that all of the DS3000 device names and paths are correct, and that AIX recognizes the dacs and hdisks.

Using the mpio_get_config –Av command

Use the **AIX mpio_get_config -Av** command to perform the following verification tasks:

1. Correlate AIX hdisk numbers to the logical drive name displayed in the Storage Manager Client.

   Note: The logical drive name from the Storage Manager client is displayed under the “User Label” heading.

2. Make sure that logical drives are on the preferred DS3400 controller (controller A or B).
3. Make sure that the correct number of storage controllers are discovered:

   Controller count: 1 (Indicates a single-controller configuration)
   Controller count: 2 (Indicates a dual-controller configuration)

4. Make sure that the “Partition count” matches the number of storage partitions configured to the system on the Storage Manager client.

   The following example shows the output of the **mpio_get_config –Av** command for a dual-controller DS3400 storage subsystem:

```
# mpio_get_config -Av
Frame id 0:
    Storage Subsystem worldwide name: 60ab800374244000047064e0
    Controller count: 2
    Partition count: 2
    Partition 0:
    Storage Subsystem Name = 'Accounting_FEB'
        hdisk      LUN #   Ownership          User Label
        hdisk35        0   B (preferred)       E-1-S5
        hdisk36        1   B (preferred)       F-1-S5
        hdisk37        2   B (preferred)       G-1-S5
        hdisk38        3   B (preferred)       H-1-S5
        hdisk39        4   A (preferred)       E-2-S5
        hdisk40        5   A (non-preferred)   F-2-S5
```

Using the lspath command

Use the **lspath** command to verify the following information for each DS3400 hdisk:

1. Make sure that the expected number of paths are detected.
2. Verify that the path status is Enabled.

   The following example shows the output of the **lspath** command for a dual-controller DS3400 storage subsystem:

```
# lspath|sort

Enabled hdisk10  fscsi1
Enabled hdisk10  fscsi1
Enabled hdisk10  fscsi3
Enabled hdisk10  fscsi3
Enabled hdisk11  fscsi1
Enabled hdisk11  fscsi1
Enabled hdisk11  fscsi3
Enabled hdisk11  fscsi3
Enabled hdisk12  fscsi1
Enabled hdisk12  fscsi1
Enabled hdisk12  fscsi3
Enabled hdisk12  fscsi3
```

The **lspath** command example shows that all paths are in the Enabled state and that each hdisk has four paths.

Using the lsdev command

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

If you are using in-band management, make sure that the correct number of dac devices are discovered by using the **lsdev -C|grep dac** command on the AIX system.

The following example shows the output of the **lsdev -C|grep dac** command:

```
# lsdev -C |grep dac
dac0       Available 03-00-02      DS3/4K PCM User Interface
dac1       Available 08-08-02      DS3/4K PCM User Interface
```

This **lsdev** command example shows two dac devices. One dac device should be detected for each controller in the DS3400 storage subsystem. Dac devices are detected only if the access LUN is mapped to the storage partitions for this host system.

Note: A DS3200 storage subsystem that is connected to SAS connectivity modules in a BladeCenter unit can discover only one dac device.

Each dac has its own location code, or path, which is represented by the values 03-00-02 and 08-08-02. 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 IBM System p server to identify device locations.

Using the lsattr command

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

Use the **lsattr –El** hdisk command to view the current attributes for each DS3000 associated hdisk. Make sure that the “algorithm” and “reserve_policy” attributes are set to the intended values.

The following example shows the output of the **lsattr –El hdisk30** command. In this example, the current attributes of hdisk30 are shown:

```
lsattr -El hdisk30
PCM             PCM/friend/otherapdisk           Path Control Module              False
PR_key_value    none                             16 Bit LVD SCSI Disk Drive       True
algorithm       round_robin                      Algorithm                        True
clr_q           no                               Device CLEARS its Queue on error True
cntl_delay_time 0                                Controller Delay Time            True
cntl_hcheck_int 0                                Controller Health Check Interval True
cntl_hcheck_int 0                                Controller Health Check Interval True
dist_err_pcnt   0                                Distributed Error Percentage     True
dist_tw_width   50                               Distributed Error Sample Time    True
hcheck_cmd      inquiry                          Health Check Command             True
hcheck_interval 60                               Health Check Interval            True
hcheck_mode     nonactive                        Health Check Mode                True
location                                         Location Label                   True
lun_id          0x7000000000000                  Logical Unit Number ID           False
max_transfer    0x40000                          Maximum TRANSFER Size            True
node_name       0x200800a0b824581a               FC Node Name                     False
pvid            000c27fe9a2184350000000000000000 Physical volume identifier       False
q_err           yes                              Use QERR bit                     True
q_type          simple                           Queuing TYPE                     True
queue_depth     4                                Queue DEPTH                      True
reassign_to     120                              REASSIGN time out value          True
reserve_policy  no_reserve                       Reserve Policy                   True
rw_timeout      30                               READ/WRITE time out value        True
scsi_id         0x7a0700                         SCSI ID                          False
start_timeout   60                               START unit time out value        True
ww_name         0x203800a0b824581a               FC World Wide Name               False
```

Changing the "algorithm" and "reserve_policy" attributes

To correctly use the planned configuration, the algorithm and reserve_policy attributes need to be set correctly.

Using example 1 in “Preparing for a SAN attachment (DS3400 only)” on page 25, if attribute “algorithm” is not changed to round_robin, only one of the HBAs is used until the first HBA fails. After the first HBA has failed, the second HBA will service the I/Os. If the “algorithm” is set to round_robin, the I/Os will alternate between the two HBAs.

Note: Use the **chdev** command to change the attributes.

For example, to change the attributes of the algorithm to round_robin and the reserve_policy to no_reserve for hdisk30, the command is:

```
#  chdev -l hdisk30 -a algorithm=round_robin –a reserve_policy=no_reserve
```

Viewing and setting Object Data Manager (ODM) attributes

Some of the Object Data Manager (ODM) attributes are for information purposes only. These information-only attributes show how the DS3000 storage subsystem is configured or its current state. You can modify other attributes using SMIT or by using the AIX **chdev -p** command.

Using the lsattr command to view ODM attributes

To view the Object Data Manager (ODM) attribute settings for fcs, fscsi, dacs, and hdisks, use the **lsattr** command to perform the following tasks::

- To view the default settings, type `lsattr -Dl`.

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

The following example displays the attributes for the fcs device.

```
# lsattr -El fcs0
bus_intr_lvl  305        Bus interrupt level                                 False
bus_io_addr   0xff800    Bus I/O address                                     False
bus_mem_addr  0xfff7e000 Bus memory address                                  False
init_link     al         INIT Link flags                                     True
intr_priority 3          Interrupt priority                                  False
lg_term_dma   0x800000   Long term DMA                                       True
max_xfer_size 0x100000   Maximum Transfer Size                               True
num_cmd_elems 200        Maximum number of COMMANDS to queue to the adapter True
pref_alpa     0x1        Preferred AL_PA                                     True
sw_fc_class   2          FC Class for Fabric                                 True
```

The following example displays the attributes for the fscsi device.

```
# lsattr -El fscsi0
attach       switch       How this adapter is CONNECTED         False
dyntrk       no           Dynamic Tracking of FC Devices        True
fc_err_recov delayed_fail FC Fabric Event Error RECOVERY Policy True
scsi_id      0x7c0e00     Adapter SCSI ID                       False
sw_fc_class  3            FC Class for Fabric                   True
```

The following example displays the attributes for the dac device.

```
# lsattr -El dac0
PCM       PCM/friend/ds4k-ui Path Control Module    False
lun_id    0x1f000000000000   Logical Unit Number ID False
node_name 0x200400a0b83743b1 FC Node Name           False
scsi_id   0xef               SCSI ID                False
ww_name   0x202400a0b83743b1 FC World Wide Name     False
```

The following example displays the attributes for the hdisk device

```
# lsattr -El hdisk30

PCM             PCM/friend/otherapdisk Path Control Module             False
PR_key_value    none                   16 Bit LVD SCSI Disk Drive      True
algorithm       fail_over              Algorithm                       True
autorecovery    yes                    N/A                             True
clr_q           no                     Device CLEARS its Queue on error True
cntl_delay_time 0                      Controller Delay Time           True
cntl_hcheck_int 0                      Controller Health Check Interval True
cntl_hcheck_int 0                      Controller Health Check Interval True
cntl_hcheck_int 0                      Controller Health Check Interval True
cntl_hcheck_int 0                      Controller Health Check Interval True
dist_err_pcnt   0                      Distributed Error Percentage    True
dist_tw_width   50                     Distributed Error Sample Time   True
hcheck_cmd      inquiry                Health Check Command            True
hcheck_interval 60                     Health Check Interval           True
hcheck_mode     nonactive              Health Check Mode               True
location                               Location Label                  True
lun_id          0x11000000000000       Logical Unit Number ID          False
max_transfer    0x40000                Maximum TRANSFER Size           True
node_name       0x200400a0b824588d     FC Node Name                    False
pvid            none                   Physical volume identifier      False
q_err           yes                    Use QERR bit                    True
q_type          simple                 Queuing TYPE                    True
queue_depth     10                     Queue DEPTH                     True
reassign_to     120                    REASSIGN time out value         True
reserve_policy  single_path            Reserve Policy                  True
rw_timeout      30                     READ/WRITE time out value       True
scsi_id         0x11100                SCSI ID                         False
start_timeout   60                     START unit time out value       True
ww_name         0x202400a0b824588d     FC World Wide Name              False
```

Setting the queue depth for hdisk devices

Setting the *queue_depth* attribute to the applicable value is important for system performance. If you have a large DS3400 configuration with many logical drives and hosts attached, use this setting for high performance.

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

Calculating the maximum queue depth: For DS3400 storage subsystems, 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, has a maximum queue depth of 16:

2048 / (4 × 32) = 16

Setting the 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
```

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.

Viewing the capacity of hdisk

The **bosinfo -s <hdisk#>** command can be used to display the current capacity of the hdisk. By default, the capacity is displayed in megabytes. The **bootinfo -s** command displays the updated capacity after a Dynamic Volume Expansion operation.

Additional AIX configuration information

The following sections describe additional AIX configuration information.

Using fast I/O failure for Fibre Channel devices

I/O failover occurs when a Fibre Channel adapter device driver detects a link failure on the fabric between the switch and the DS3000 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 device driver detects that a link was lost between the switch and the DS3000 storage subsystem, it waits a few seconds to enable the fabric to stabilize. If **fast_fail** is specified, the adapter then begins failing all I/O at the adapter device 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 enable 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 DS3000 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 DS3000 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

AIX dynamically tracks Fibre Channel devices when they are moved from one Fibre Channel port on a switch to another port. This enables the dynamic movement of a Fibre Channel connection by suspending I/O for 15 seconds while the move occurs. If a cable move is completed in less than 15 seconds, this feature prevents the logical drives from being moved off the preferred path.

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 intended 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 DS3000 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 DS3000 storage subsystem, but it requires manual intervention for AIX to recognize the new logical drive capacity. This section explains how to use DVE with AIX.

Before you begin: Make sure that there is sufficient available free capacity within the array. Free capacity can be determined using the DS3000 Storage Manager software. In the **Summary/Arrays & Logical Drives** view of the Subsystem

Management window, expand the desired array and scroll to the bottom of the logical drives list. Available free capacity is an entry in the logical drives list.

If there is insufficient free capacity and extra drives are available, perform a dynamic capacity expansion (DCE) operation before performing the DVE operation. A DCE operation increases the capacity of the array by adding unused physical disks that are already installed in the drive enclosures.

Performing a dynamic capacity expansion operation

To perform a dynamic capacity expansion operation, complete the following steps:

1. In the Modify/Add Free Capacity (Drives) view, select the array whose capacity is to be increased.
2. Select **Next**.
3. Select an available drive in the “Add capacity to array” window.
4. Select **Finish**.

Note: The available free disks and their capacity are displayed in the window.

After the operation is started, a clock is displayed beside every logical drive in the array and remains until the operation is completed. The process might take several hours and must be allowed to complete before AIX intervention.

Note: The time required for the operation to be completed depends on the I/O activity, capacity being added, drive technology, and other factors.

Performing a dynamic volume expansion operation

Performing a dynamic volume expansion requires the functionality of the Storage Manager software SMcli interface or the Storage Manager software script editor. The following example shows the command syntax that is required to run the dynamic volume expansion from the script editor window on a logical drive named Engineering_FEB. The capacity added is 1 GB.

```
set logicalDrive ["Engineering_FEB"] addCapacity=1GB;
```

Note: After the operation is completed, you can view the new capacity on the AIX system by using the **bootinfo –s hdiskX** command.

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

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

   ```
   umount mount_point
   ```

   where *mount_point* is the name of the file system 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 message that the logical drive group cannot be imported to AIX 5.1 or earlier, ignore the warning message. The error message is incorrect.
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 the logical drive while the logical drive group is activated in classic or enhanced concurrent mode.
3. You cannot resize the root logical drive group.

Using the AIX autorecovery feature

AIX supports host-initiated recovery of logical drives to the preferred path after a fault is corrected. You can enable the autorecovery characteristic by setting the hdisk attribute Autorecovery to **yes**. The default value for the Autorecovery attribute for hdisks is **No**.

To enable autorecovery for hdisk20, type the following command.

```
# chdev -l hdisk20 -a autorecovery=yes
```

Note: For the autorecovery mechanism to work correctly, the logical drives must be located on their preferred path before a path fault occurs.

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 adapter replacement
- DS3000 Storage Manager software
- Fibre Channel zoning procedures (required if a switch is used to interconnect DS3000 storage subsystems and the AIX host, and zoning is based on WWPN)

Attention: If the procedures are not followed as they are documented in this section, loss of data availability can occur. Make sure that you read and understand all of the requirements and steps in this section before you begin the HBA hot-swap procedure.

Known issues and restrictions

Before you perform a hot-swap operation, read the following known issues and restrictions:

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

- You must replace a defective HBA with the *same* model HBA, and install it in the *same* PCI slot.

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

  Important: No other variations of replacement scenarios are currently supported.
- Hot-swapping an HBA is not supported if the replacement of the HBA will remove the last viable path to any hdisk.

Collecting system data

To collect data from the system, complete the following steps:

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:

   ```
   # lsdev -C |grep dac
   dac0       Available 00-08-01      DS3/4K PCM User Interface
   dac1       Available 00-08-01      DS3/4K PCM User Interface
   ```

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 is 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
   ```

Replacing the hot-swap HBA

After you complete the procedures in “Collecting system data” on page 40, 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 to hot-swap the HBA 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 the targeted HBA. A window opens 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 steps 1 to 4 in this procedure are completed successfully, make sure that you have the following results:
   - The defective HBA is removed from the system.
   - The replacement Fibre Channel HBA is powered on.
   - The associated fcs*x* device is in the Defined state.
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 applicable 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 40.
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 is used in the next procedure to manually map the replacement HBA WWPN to the storage subsystems.
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 DS3000 storage subsystem” on page 42.

Mapping the new WWPN to the DS3000 storage subsystem

For each DS3000 storage subsystem that is affected by the hot-swap procedure, to map the worldwide port name (WWPN) of the new HBA to the storage subsystem, complete the following steps:

1. Start the DS3000 Storage Manager software and open the Subsystem Management window.
2. In the Subsystem Management window, select the **Modify** tab.
3. Locate the defective HBA, select **HBA**, and then click **Replace HBA**.
4. Locate the new HBA WWPN from the list, modify the alias, if you want, and then click **OK** to activate the change.

Completing the HBA hot-swap procedure

To finish replacing the hot-swap HBA, complete the following steps:

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

   Important: Skip this step if the HBA is directly attached to the DS3000 storage subsystem, or if the Fibre Channel switch zoning is based on port numbers instead of WWPNs. If you need to modify the zoning, failure to correctly do so prevents the HBA from accessing the storage subsystem.
3. Run the **cfgmgr** command.
4. Use the **lsdev –C** command to verify that the replaced fcs*x* device and its associated dacs are placed in the Available state. See the output from the following example.

```
# lsdev -C |grep fcs
fcs0       Available 00-08         FC Adapter
fcs1       Available 00-09         FC Adapter
fcs2       Available 05-08         FC Adapter
fcs3       Available 05-09         FC Adapter
```

5. Using the Storage Manager Client, manually redistribute the logical drives to the preferred paths unless the autorecovery attribute for the hdisks is set to Yes.

   The autorecovery functionality starts redistributing drives that are not on the preferred path after the link health is detected as Healthy for 5 minutes.

   To manually redistribute the logical drives, complete the following steps:

   a. Click the **Support** tab.

   b. Select **Manage Controllers**, and then select **Redistribute Logical Drives**.
6. Make sure that the disks stay on the preferred path by using one or both of the following methods:

   Using the AIX system

   Run the **mpio_get_config -Av** command and make sure that the drives are on the expected path.

   Using the Storage Manager software

   In the Enterprise Management window, make sure that the storage subsystems are Optimal. If they are not Optimal, make sure that any

drives that are part of the storage subsystems that are involved with hot-swap process are *not* listed in the Recovery GURU.

7. Using the AIX system, run the **lspath** command to make sure that all paths to the hdisks are Enabled.

DS3000 error log messages

This section describes the possible errors that might be reported in the AIX error log. You can view the AIX error log by running the **errpt -a** command.

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

Note: The snapshot references in the following list equate to FlashCopy when reported on a DS3000 storage subsystem.

- **SC_DISK_PCM_ERR1:** *Subsystem Component Failure*

  The storage subsystem has returned an error indicating that some component (hardware or software) of the storage subsystem has failed. The detailed sense data identifies the failing component and the recovery action that is required. Failing hardware components should also be shown in the Storage Manager software, so the placement of these errors in the error log is advisory and is an aid for your technical-support representative.

- **SC_DISK_PCM_ERR2:** *Array Active Controller Switch*

  The active controller for one or more hdisks associated with the storage subsystem has changed. This is in response to some direct action by the AIX host (failover or autorecovery). This message is associated with either a set of failure conditions causing a failover or, after a successful failover, with the recovery of paths to the preferred controller on hdisks with the autorecovery attribute set to Yes.

- **SC_DISK_PCM_ERR3:** *Array Controller Switch Failure*

  An attempt to switch active controllers has failed. This leaves one or more paths with no working path to a controller. The AIX MPIO PCM will retry this error several times in an attempt to find a successful path to a controller.

- **SC_DISK_PCM_ERR4:** *Array Configuration Changed*

  The active controller for an hdisk has changed, usually due to an action not initiated by this host. This might be another host initiating failover or recovery, for shared LUNs, a Redistribute operation from the Storage Manager software, a change to the preferred path in the Storage Manager software, a controller being taken offline, or any other action that causes the active controller ownership to change.

- **SC_DISK_PCM_ERR5:** *Array Cache Battery Drained*

  The storage subsystem cache battery has drained. Any data remaining in the cache is dumped, and is vulnerable to data loss until it is dumped. Caching is not normally allowed with drained batteries, unless the administrator takes action to enable it within the Storage Manager software.

- **SC_DISK_PCM_ERR6:** *Array Cache Battery Charge Is Low*

  The storage subsystem cache batteries are low and need to be charged or replaced.

- **SC_DISK_PCM_ERR7:** *Cache Mirroring Disabled*

  Cache Mirroring is disabled on the affected hdisks. Normally, any cached write data is kept within the cache of both controllers, so that if either controller fails,

there is still a good copy of the data. This is a warning message stating that loss of a single controller will result in data loss.

- **SC_DISK_PCM_ERR8:** *Path Has Failed*

  The I/O path to a controller has failed or gone offline.
- **SC_DISK_PCM_ERR9:** *Path Has Recovered.*

  The I/O path to a controller has resumed and is back online.
- **SC_DISK_PCM_ERR10:** *Array Drive Failure*

  A physical drive in the storage array has failed and should be replaced.
- **SC_DISK_PCM_ERR11:** *Reservation Conflict*

  A PCM operation has failed due to a reservation conflict. This error is not currently issued.
- **SC_DISK_PCM_ERR12:** *Snapshot Volume's Repository Is Full*

  The snapshot volume repository is full. Write actions to the snapshot volume will fail until the repository problems are fixed.
- **SC_DISK_PCM_ERR13:** *Snapshot Op Stopped By Administrator*

  The administrator has halted a snapshot operation.
- **SC_DISK_PCM_ERR14:** *Snapshot repository metadata error*

  The storage subsystem has reported that there is a problem with snapshot metadata.
- **SC_DISK_PCM_ERR15:** *Illegal I/O - Remote Volume Mirroring*

  The I/O is directed to an illegal target that is part of a Remote Volume Mirroring pair (the target volume rather than the source volume).
- **SC_DISK_PCM_ERR16:** *Snapshot Operation Not Allowed*

  A snapshot operation that is not allowed has been attempted.
- **SC_DISK_PCM_ERR17:** *Snapshot Volume's Repository Is Full*

  The snapshot volume repository is full. Write actions to the snapshot volume will fail until the repository problems are fixed.
- **SC_DISK_PCM_ERR18:** *Write Protected*

  The hdisk is write-protected. This can happen if a snapshot volume repository is full.
- **SC_DISK_PCM_ERR19:** *Single Controller Restarted*

  The I/O to a single-controller storage subsystem is resumed.
- **SC_DISK_PCM_ERR20:** *Single Controller Restart Failure*

  The I/O to a single-controller storage subsystem is not resumed. The AIX MPIO PCM will continue to attempt to restart the I/O to the storage subsystem.

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, do not automatically redirect logical drives to the preferred paths.

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

1. Repair or replace any faulty components. For more information, see the *Installation, User's, and Maintenance Guide* for the applicable DS3000 storage subsystem.

2. To redistribute logical drives to their preferred paths, on the Subsystem Management window, selecting the **Support** tab. Select **Manage controller**, and then select **Redistribute Logical Drives**.

   Note: If a large number of logical drives are configured on the DS3000 storage subsystem, redistributing the logical drives might take 2 hours or more to be completed, depending on the system activity level.

3. On the AIX systems, run the **mpio_get_config -Av** command to make sure that all the logical drives are on the preferred controller, as shown in the following example.

```
# mpio_get_config -Av
 Frame id 3:
     Storage Subsystem worldwide name: 60ab80024585d000047fdb68
     Controller count: 2
     Partition count: 1
     Partition 0:
     Storage Subsystem Name = 'C4'
         hdisk      LUN #   Ownership          User Label
         hdisk53        0   B (preferred)      A-1-S5
         hdisk54        1   A (preferred)      A-10-S1
         hdisk55        2   B (preferred)      A-11-S0
         hdisk56        3   A (preferred)      A-12-S0
         hdisk57        4   A (preferred)      A-2-S5
         hdisk58        5   B (preferred)      A-3-S5
         hdisk59        6   A (preferred)      A-4-S5
```

Logical drives that are not on the preferred path will display “(non-preferred)” under the Ownership heading.

Chapter 4. Installing and configuring the Storage Manager software on POWER-based Linux host systems

You can use the DS3000 Storage Manager software with the following Linux distributions for POWER-based host operating systems:

- SUSE Linux Enterprise Server 9 (SLES 9)
- SUSE Linux Enterprise Server 10 (SLES 10)
- Red Hat Enterprise Linux 4 (RHEL 4)
- Red Hat Enterprise Linux 5 (RHEL 5)

Note: Before you begin any of the tasks in this chapter, see "SUSE Linux Enterprise Server system requirements" and "Red Hat Enterprise Linux system requirements" on page 49 for the system requirements for your Linux distribution.

For more information about Linux on POWER and IBM System p support, see the following Web sites:

Linux on POWER resource center
http://www.ibm.com/servers/enable/linux/power/

Linux on System p support
http://www.ibm.com/servers/eserver/pseries/linux/

Connecting host bus adapters in a Fibre Channel switch environment

When you connect Fibre Channel host bus adapters in host servers to DS3000 storage subsystem host ports in a Fibre Channel switch environment, you should establish one-to-one zones between the host bus adapter and DS3000 host ports. If multiple adapters or controller ports are in the same zone, you might encounter extended failover delays or additional failovers because device resets are being propagated throughout the zone.

For instructions for setting up a zone inside a Fibre Channel switch, see the documentation that comes with the Fibre Channel switch.

You can find links to switch documentation at http://www.ibm.com/servers/storage/support/san/index.html.

SUSE Linux Enterprise Server system requirements

This section lists the minimum hardware and software requirements for a POWER-based host running SLES 9 or SLES 10 attached to a DS3000 storage subsystem that is managed by the Storage Manager software.

Important: See the Storage Manager software readme file for Linux on POWER for information about UTM LUN limitations. To get to the readme file, see "Obtaining the documentation from the IBM Support Web site" on page 1.

Host hardware requirements

The following minimum hardware is required.

Note: For the most current host support information, see the latest DS3000 interoperability matrix.

- **DS3200 support:**
  - IBM BladeCenter H chassis or IBM BladeCenter E chassis with two IBM BladeCenter SAS Connectivity Modules
  - IBM BladeCenter JS12, IBM BladeCenter JS21, and IBM BladeCenter JS22, using the IBM BladeCenter SAS Expansion Card (CFFv)
  - One BladeCenter unit can connect to only one DS3200 storage subsystem.

    Note: To configure a BladeCenter unit and a DS3200 storage subsystem, connect host port 1 on a DS3200 controller to a BladeCenter SAS Connectivity Module in a BladeCenter unit. Connect the other host port 1 on the second DS3200 controller to a second BladeCenter SAS Connectivity Module in a BladeCenter unit. No other external devices can be connected to the SAS connectivity modules.
- **DS3300 support:**

  IBM BladeCenter JS21 and IBM BladeCenter JS22, using the QLogic iSCSI Expansion Card for IBM @server BladeCenter
- **DS3400 support:**
  - Any IBM System p or POWER server using the IBM POWER4, POWER5, or POWER6 64-bit PowerPC architecture that supports one of the following Fibre Channel host bus adapters:
    - FC 6239
    - FC 5716 / 1957 / 1977
    - FC 5758 / 1905
    - FC 5759 / 1910
    - FC 5773
    - FC 5774
  - IBM BladeCenter JS21 and IBM BladeCenter JS22, using the IBM @server BladeCenter Fibre Channel expansion cards
- **Management station for the Storage Manager software**

  You can use the SLES host as the management station for the Storage Manager software.
  - For Microsoft® Windows® or i386-based Linux management stations, see the minimum requirements listed for those operating systems in the *IBM System Storage DS3000 Storage Manager Version 10 Installation and Support Guide for Microsoft Windows Server, Linux, Novell NetWare, and VMware ESX Server*.
  - For AIX or RHEL management stations, see the minimum requirements listed in the applicable chapters in this document.

Software requirements

For the latest supported versions of the SUSE Linux operating system and supported versions of host bus adapter device drivers, see the Storage Manager software readme file for Linux on POWER.

Note: To get to the readme file, see “Obtaining the documentation from the IBM Support Web site” on page 1.

Red Hat Enterprise Linux system requirements

This section lists the minimum hardware and software requirements for a POWER-based host running RHEL 4 or RHEL 5 attached to a DS3000 storage subsystem that is managed by the Storage Manager software.

Important: See the Storage Manager software readme file for Linux on POWER for information about UTM LUN limitations. To get to the readme file, see “Obtaining the documentation from the IBM Support Web site” on page 1.

Host hardware requirements

The following minimum hardware is required:

- **DS3200 support:**
  - Supported BladeCenter unit: IBM BladeCenter H chassis or IBM BladeCenter E chassis with two IBM BladeCenter SAS Connectivity Modules
  - Supported blade servers: IBM BladeCenter JS12, IBM BladeCenter JS21, and IBM BladeCenter JS22, using the IBM BladeCenter SAS Expansion Card (CFFv)
  - One BladeCenter unit can connect to only one DS3200 storage subsystem.

    Note: To configure a BladeCenter unit and a DS3200 storage subsystem, connect host port 1 on a DS3200 controller to a BladeCenter SAS Connectivity Module in a BladeCenter unit. Connect the other host port 1 on the second DS3200 controller to a second BladeCenter SAS Connectivity Module in a BladeCenter unit. No other external devices can be connected to the SAS connectivity modules.
- **DS3300 support:**

  IBM BladeCenter JS21 and IBM BladeCenter JS22, using the QLogic iSCSI Expansion Card for IBM @server BladeCenter
- **DS3400 support:**
  - Any IBM System p or POWER server using the IBM POWER4, POWER5, or POWER6 64-bit PowerPC architecture that supports one of the following Fibre Channel host bus adapters:
    - FC 6239
    - FC 5716 / 1957 / 1977
    - FC 5758 / 1905
    - FC 5759 / 1910
    - FC 5773
    - FC 5774
  - IBM BladeCenter JS21 and IBM BladeCenter JS22, using the IBM @server BladeCenter Fibre Channel expansion cards
- **Management station for the Storage Manager software**

  You can use the RHEL host as the management station for the Storage Manager software.
  - For Windows or i386-based Linux management stations, see the minimum requirements listed for those operating systems in the *IBM System Storage*

DS3000 Storage Manager Version 10 Installation and Support Guide for Microsoft Windows Server, Linux, Novell NetWare, and VMware ESX Server.

– For AIX or RHEL management stations, see the minimum requirements listed in the applicable chapters in this document.

Software requirements

For the latest supported versions for the Red Hat Enterprise Linux operating system and supported versions of host bus adapter device drivers, see the Storage Manager software readme file for Linux on POWER.

Note: To get to the readme file, see "Obtaining the documentation from the IBM Support Web site" on page 1.

Linux host restrictions

Before you begin working with the Storage Manager software, read the following restrictions. These restrictions apply to both the SLES and RHEL operating systems, except as otherwise noted.

Important: See the Storage Manager software readme file for Linux on POWER for information about UTM LUN limitations. To get to the readme file, see "Obtaining the documentation from the IBM Support Web site" on page 1.

FC SAN and connectivity restrictions

- Dual-path configurations are supported only with the IBM RDAC failover driver.
- Single-switch configurations are allowed, but each HBA and DS3000 controller combination must be in a separate SAN zone or VLAN.
- For Fibre Channel, other storage devices, such as tape devices or other disk storage, must be connected through separate HBAs and SAN zones.
- As of the date of this document, clustering is not supported.

BladeCenter unit and DS3200 connectivity restrictions:

AIX supports the attachment of only one dual-controller DS3200 to a BladeCenter unit.

The attachment requires two SAS connectivity modules. One SAS connectivity module is connected by a single cable to host port 1 of controller A. The second SAS connectivity module is connected by a single cable to host port 1 of controller B. No other external devices can be connected to the SAS connectivity modules or to host ports 2 or 3 of the DS3200 controller modules.

Linux RDAC restrictions (AVT/ADT)

This Linux RDAC release does not support auto-volume transfer/auto-disk transfer (AVT/ADT) mode. AVT/ADT is disabled by default in the Linux storage partitioning host type. Use the Linux storage partitioning host type for Linux partitions.

Partitioning restrictions

- Dual-path configurations are supported only with the IBM RDAC failover driver.
- 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 a LUN between 0 and 31 that is not a UTM or access logical drive.

Installing the Storage Manager software

This section describes the procedures for installing the Storage Manager software.

Preparing for the software installation

Before you install the DS3000 Storage Manager software, you might need the following items to prepare for the installation:

- For the Linux on POWER operating system, make sure that you are running the operating-system-required maintenance level. (See the readme file for the required maintenance levels and any additional fixes.)
- HBA drivers
- The most up-to-date controller firmware for your particular storage subsystem as listed in “SUSE Linux Enterprise Server system requirements” on page 47 or “Red Hat Enterprise Linux system requirements” on page 49.
- The IP addresses for RAID controllers (for out-of-band management only)
- Additional documentation for switches and HBAs, if needed
- Applicable host software kit. The host software kit grants you permission to attach host servers using the applicable operating system for the DS3000. The kit includes the most up-to-date DS3000 software and controller firmware for your particular storage subsystem.
- Applicable host attachment support CD for your operating system, which includes the most up-to-date DS3000 software and controller firmware for your particular storage subsystem.

For the latest controller firmware, see http://www.ibm.com/servers/storage/support/disk/.

Best practice: When you start the host in preparation for installation of the Storage Manager software packages, make sure that DS3000 devices are not yet attached to the host system. Attaching DS3000 devices before you configure the host increases system startup time, and can complicate the installation process. Attach the DS3000 devices after you complete the procedures in this section, and before you install the Storage Manager software.

Failover drivers

A Linux on POWER host system requires either the RHEL or the SLES Redundant Disk Array Controller (RDAC) failover driver for I/O path redundancy. The failover driver monitors I/O paths. If a component failure occurs in one of the I/O paths, the failover driver reroutes all I/O to another path.

For information about installing the RDAC multipath driver, see “Installing the RDAC multipath proxy driver” on page 58.

Configuring the Linux host

After the Linux operating system is installed and before you install the Storage Manager software, complete the procedures that are described in this section to configure your Linux host for the DS3000 storage subsystem.

Installing the HBA device driver and utility packages

The type of HBA device driver that you need depends on your configuration. There are different drivers for Fibre Channel, SAS, and iSCSI host bus adapters. There might also be different device drivers for the SLES and RHEL operating systems.

Note: For the latest supported versions of the following packages, see the DS3000 Storage Manager readme file for Linux on POWER.

SLES 10 with SP1 or RHEL 5 HBA driver support

For systems using SLES 10 with SP1 or RHEL 5, unless otherwise noted, use the HBA device driver that comes with the operating system.

SLES 9 and RHEL 4 HBA driver support

You can download the following items from the DS3000 support Web site:

- **Driver package**

  Compiles changes in the configuration file (/etc/lpfc.conf)
- **Application package**

  Monitors the HBAs and facilitates the downloading of the HBA firmware.
- **Release notes**

  Contains the instructions for upgrading the device driver.

To install the HBA driver, complete the following steps:

1. Create a directory for the HBA driver package.
2. Go to http://www.ibm.com/servers/storage/support/disk/ and download the applicable driver package to the directory that you just created.
3. Uncompress the file by typing the following command:

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

   Note: If the RDAC is already installed on the system, the RDAC driver must be unloaded or removed before installing the HBA device driver. Do not load or unload the RDAC driver stack, which includes the mpp_Upper and mpp_Vhba tools and the low level host bus adapter device driver, by using the modprobe kernel module utility. Any use of the modprobe utility on the RDAC driver stack is not supported. Using rmmod to remove all the drivers in the RDAC driver stack, in order, is also not supported. Each time you must unload the driver stack, restart the system.
4. To update a physical HBA driver, complete the following steps:
   a. Get the latest HBA drivers that are supported by MPP.
   b. Install the drivers.
   c. After the driver is installed, change back to the initial Linux RDAC installation directory, and reinstall the RDAC driver using the following commands:

      ```
      # make clean
      # make uninstall
      # make
      # make install
      ```
5. Restart the server.

Installing the optional application kit

To install the optional application kit, complete the following steps:

1. Create a directory for the application package.
2. Go to http://www.ibm.com/servers/storage/support/disk/ and download the applicable driver package to the directory that you just created.
3. Uncompress the file by typing the following command:

   ```
   # tar -zxvf package_name.tgz
   ```
4. In the directory, find the readme.txt file, and follow the instructions for installing the driver package.
5. Restart the server.

Recording the HBA worldwide port names

During the installation process, you need to know the HBA worldwide port names. To view and record the Fibre Channel HBA WWPNs or iSCSI HBA IQN, complete the following steps for either Emulex-based HBAs or QLogic-based HBAs.

For Emulex-based HBAs:

1. Type the following command to view a the HBAs:

   ```
   # ls /proc/scsi/lpfc
   ```

   A list containing the assigned numbers for all of the HBAs that are discovered on the system is displayed.
2. In the list, identify the host number of the HBA that is connected to the DS3000 storage subsystem and go to that directory by typing the following command:

   ```
   # cd /sys/class/scsi_host/host#
   ```

   where *#* is the assigned HBA number that you discovered in step 1.
3. Type the following command to view the WWPN of the HBA and connected DS3000 controllers:

   ```
   # cat port_name
   ```

   where *port_name* is the assigned HBA number. The *port_name* value is equivalent to the WWPN.

For QLogic-based HBAs:

1. Type the following command to view a the HBAs:

   ```
   # ls /proc/scsi/qla2*

   or

   # ls /proc/scsi/qla4*
   ```

   A list containing the assigned numbers for all of the HBAs that are discovered on the system is displayed.
2. For each HBA that is connected to the DS3000 storage subsystem, type the following command:

   ```
   # cat /proc/scsi/qla2300/x
   ```

   where *x* is the assigned HBA number that you discovered in step 1.

   The following information is displayed:

   ```
   scsi-qla#-adapter-port-xxxxxxxxxxxxxx
   ```

The adapter port value *xxxxxxxxxxxxxx* provides the Fibre Channel WWPN or iSCSI IQN.

Storage Manager installation and configuration

After the host configuration procedure is completed, install and configure the Storage Manager software by following the procedures in this section.

Note: Before you begin the Storage Manager software installation, make sure that the DS3000 devices are attached to the host.

You can install all of the Storage Manager software packages automatically by using the Storage Manager software installation wizard or you can install the individual packages by using silent mode installation.

Installing Storage Manager using the installation wizard

The Storage Manager software installation wizard is a Java-based interactive method for choosing which packages to automatically install on your host system. The installation wizard installs the following software packages on the host:

- SMruntime
- SMclient
- SMagent
- SMutil

Requirements:

- DS3000 Storage Manager version 02.70, or later
- xservices must be operational

If you are installing Storage Manager version 02.70, if the 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 standalone host software packages using the procedures described in “Installing the Storage Manager software in console mode” on page 55.

Note: If you use the wizard to install the Storage Manager software, you might need to increase the size of some partitions to accommodate the installation.

To install the DS3000 Storage Manager software using the installation wizard, complete the following steps. Adjust the steps as necessary for your specific installation.

1. Copy the Storage Manager software SMIA package file from the *System Storage DS3000 Support software for Linux on POWER* CD, or download the package file from the DS3000 support Web site, to a directory on your system.
2. Open the file on your system by typing the following command:

   ```
   # sh SMIA-LINUXPPC-02.70.A5.xx.bin
   ```

   The Storage Manager Installation wizard 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 the Storage Manager software packages
   - Management Station—Installs SMruntime and SMclient
   - Host—Installs SMruntime, SMagent, and SMutil

- Custom—Enables you to select which packages you want to install

The DS3000 Storage Manager software is installed on your system.

Note: During the installation, you will see the question Automatically Start Monitor? This refers to the Event Monitor service. If you want to enable automatic ESM firmware synchronization, the Event Monitor must be enabled. To enable the Event Monitor, select **Automatically Start Monitor**.

Installing the Storage Manager software in console mode

Use the following steps to install the Storage Manager software with out the graphical interface.

1. From a command prompt, run the Storage Manager installation package while using the - i console parameter to forces the package to install in a non-graphical mode. For Example, in Linux, the command is:

   ```
   # sh SMIA-LINUXPPC-02.70.A5.15.bin – i console
   ```

2. Select the location (the default value is English) and press Enter.
3. Read the End User License Agreement (EULA). To continue, press **Y** (for Yes) to accept the EULA.
4. Choose the installation type:
   - Typical
   - Management Station
   - Host
   - Customize

   Press Enter.
5. Review the Pre-Installation Summary, and then press Enter.
6. The installation is completed. Press Enter to exit the installer.

Configuring storage subsystems

If you have purchased a Storage Partitioning premium feature, make sure that the premium feature is enabled. For more information, see Chapter 7, “Enabling and using premium features,” on page 109.

To configure the Storage Manager software for a Linux on POWER system, complete the procedures in the following sections. You can configure the Storage Manager software from the Storage Manager Client software running on either a Linux on POWER or a non-Linux on POWER system.

Adding storage subsystems to the Storage Manager Client

To add storage subsystems to the Storage Manager Client (SMclient), complete the following steps:

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

   Note: For information about assigning IP addresses to the controllers, see “Step 3: Assigning IP addresses to hosts and controllers” on page 19.

2. After the storage subsystem is configured on the network, start SMclient on the host server by selecting the Storage Manager icon or by typing the following command:

   ```
   # /opt/IBM_DS3000/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 is displayed in the Enterprise Management window.

      Note: Failure to add both controllers in a dual-controller storage subsystem results in a partially-managed system, where all functions might not be available.

      The name of the storage subsystem is displayed in the Enterprise Management window.

Upgrading controller firmware and NVSRAM

For detailed procedures for downloading the firmware updates, see “Downloading controller, NVSRAM, ESM, and hard disk drive firmware” on page 94.

Creating storage partitions

To create storage partitions, complete the procedures in the following sections.

Storage partitioning overview

Before you create the storage partitions using the procedures in this section, read the following information:

- The procedures in this section assume that you have already created a physical connection between the host and the storage subsystem controllers, and that you have also connected and zoned the SAN switch (if any). If you have not completed these connections, the Storage Manager software is *not* able to list the WWPNs or IQNs of the HBAs during these procedures. In this case you will have to type the WWPNs or IQNs in the applicable fields in the procedure for defining a host and host ports.
- Create the Linux on POWER host group at the storage subsystem level. Do not create host groups at the default group level.
- Auto volume transfer/auto-disk transfer (AVT/ADT) is managed by the Linux RDAC driver. To avoid contention, AVT/ADT must be disabled. By default, AVT/ADT is disabled. Selecting the **Linux** host type automatically disables AVT/ADT.

Defining the host and host port

To define the host and host ports by using the Configure Host Access wizard, complete the following steps:

1. Select the **Configure** tab on the Subsystem Management window and select one of the following methods:
   - Create Host Access (Automatic)
   - Create Host Access (Manual)

2. Select **Create Host Access (Automatic)** if you have added a new host that is also running the host agent on the system. Otherwise go to step 5.
3. If any host are automatically configured, they are displayed in the left column as an Available Host. Select a host from the list of available hosts and click **Add**.
4. After all the available hosts are moved to the **Select Host** column, click **OK**. The host and host ports are defined and you have completed this procedure.
5. Select **Create Host Access (Manual)** to manually define the host and host port information.
6. From the Specific Host Name and Host Type window, enter the host name and then select **Linux** as the host type from the host type list. Click **Next**.
7. From the Specify HBA Host Port window, in the left pane, select the correct WWPN, SAS address, or IQN of the HBA host port that will be used to define the host. Click **Add**.

   Note: If there is not a physical connection between the host and the DS3000 controllers, the WWPNs are not displayed. In this case, you must manually enter the correct WWPN, SAS address, or IQN information by completing the following steps:

   a. Click **New**.
   b. Enter the Fibre Channel WWPN, SAS address, or iSCSI IQN information.
   c. Enter the host port alias.
   d. Click **Add**.
8. After all the host ports are added to the **Selected HBA host port** column, click **Next**.
9. From the Specify Host Group window, select one of the following options:

   No: If the host will *not* share access to the same logical drive with other hosts

   Yes: If the host will share access to the same logical drive with other hosts

   If **No** is selected, click **Next**.
   If **Yes** is selected, either enter a New Host Group name or select an existing host group from the list, then click **Next**.
10. Review the host definition details, and then click **Finish**.

The host and host ports are defined.

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.

To define a host group, complete the following steps:

1. Click the **Configure** tab on the Subsystem Management window.
2. Select **Create Host Group**.
3. Enter the new host group name (for example, Linux) and select the hosts from the left column. Click **Add** to add them to the host group.
4. When all the hosts for the host group are added, click **OK**.
5. A message is displayed indicating that the new host group is created. Click **OK**.

Mapping LUNs to a storage partition

This section contains the procedures for mapping and adding LUNs to a host or host group. These steps can be used to map LUNs to a new partition or add LUNs to an existing partition.

To map LUNs, complete the following steps:

1. Select the **Configure** tab on the Subsystem Management window.
2. Select **Create Host-to-Logical Drive Mapping**.
3. From the Select Host window, select the host or host group that the logical drives will be mapped and click **Next**.
4. From the Select Logical Drives window, select the logical drives that will be mapped to the host or host group that was selected in step 3.

   Note: If all the logical drives will be mapped to the same host or host group, then click the **Select all logical drives** check box.

   After the logical drives are selected, assign a logical unit number (LUN) from 0 to 31, to the logical drives and click **Finish**.
5. The Host-to-Logical Drive Mapping progress window is displayed. When the mappings are completed, click **OK**.
6. The mappings are completed. Select **Yes** to map additional logical drives or **No** to complete the mappings wizard.

Completing the host configuration

After you install the Storage Manager software, complete the host configuration procedures that are described in this section.

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, and that the correct HBA driver is installed.

To install RDAC, complete the following steps:

1. Download the RDAC driver package from the IBM DS3000 support Web site.

   RDAC package for SLES 9, SLES 10, and RHEL 4:
   rdac_LINUX_09.01.B5.xx_source_tar_gz

   RDAC package for SLES 10 with SP1 and RHEL 5
   rdac_LINUX_09.01.C5.xx_source_tar_gz2
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 SLES or RHEL package version number that is specified in step 1.

   A directory named **linuxrdac-***version#* (SLES 9, SLES 10, RHEL 4, and RHEL 5) is created.
4. Open the readme file that is included in the linuxrdac-*version#* directory.

5. In the readme file, find the instructions for building and installing the driver and complete the steps, which include restarting the server.

   Note: Make sure that you restart the server before you proceed to step 6.
6. Type the following command to list the installed modules:

   ```
   # lsmod
   ```
7. Make sure that the module entries are included in the **lsmod** list, as follows:
   - mppVhba
   - mppUpper
   - lpfc or qla2xxx

   Note: If the mpp_Vhba module is not displayed, the usual cause is that the server was restarted before the LUNs were assigned, so the mpp_Vhba module is not installed. If this is the case, assign the LUNs now, restart the server, and then repeat step 7.
8. Type the following command to verify the driver version:

   ```
   # mppUtil -V
   ```

9. Type the following command to verify that devices are configured with the RDAC driver:

   ```
   # ls -1R /proc/mpp
   ```

   An output similar to the following example is displayed.

```
# ls -1R /proc/mpp
/proc/mpp:
total 0
dr-xr-xr-x    4 root     root            0 Oct 24 02:56 DS3400-sys1
crwxrwxrwx    1 root     root      254,  0 Oct 24 02:56 mppVBusNode

/proc/mpp/ DS3400-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/ DS3400-sys1/controllerA:
total 0
dr-xr-xr-x    2 root     root            0 Oct 24 02:56 lpfc_h6c0t2

/proc/mpp/ DS3400-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/ DS3400-sys1/controllerB:
total 0
dr-xr-xr-x    2 root     root            0 Oct 24 02:56 lpfc_h5c0t0

/proc/mpp/ DS3400-sys1/controllerB/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
```

Note: After you install the RDAC driver, the following commands and man pages are available:

- mppUtil
- mppBusRescan
- mppUpdate
- RDAC

Performing initial device identification

The initial device identification procedures in this section apply to all the SLES and RHEL distributions, and exceptions are noted.

Complete the procedures in this section in the following order:

1. "Probing the SCSI bus"
2. "Associating each HBA WWPN with its assigned host port"
3. "Viewing the LUNs" on page 63
4. "Associating the system LUN devices with DS3000 LUNs" on page 64

Probing the SCSI bus

To probe the SCSI bus, complete one of the following procedures, according to whether RDAC is installed on your system.

- **With RDAC installed:**

  Type the following command:

  ```
  # mppBusRescan
  ```

- **Without RDAC installed:**

  Remove the modules and reload them.

  Type the following command to remove the modules:

  ```
  # rmmod hba_device_driver
  ```

  Type the following command to reload the modules:

  ```
  # modprobe hba_device_driver
  ```

Associating each HBA WWPN with its assigned host port

To associate each HBA WWPN with its assigned host port, complete the following steps:

1. Type the following commands.

   For Emulex-based HBAs:

   ```
   # cd /proc/scsi/lpfc
   # ls
   ```

   For QLogic-based HBAs:

   ```
   # cd /proc/scsi/qla2xxx
   # ls
   ```

   A list of HBAs is displayed, as shown in the following example.

   ```
   ... 5 4 3 2
   ```

2. For each HBA in the directory, type the following command:

   ```
   # cat hba#
   ```

   where *hba#* is an HBA in the directory. The HBA properties are displayed, as shown in the following example.

```
[root@x3650 /]# cat /proc/scsi/qla2xxx/2
QLogic PCI to Fibre Channel Host Adapter for QLE2462:
        Firmware version 4.00.23 [IP] , Driver version 8.01.06
ISP: ISP2432
Request Queue = 0xbee80000, Response Queue = 0xbee40000
Request Queue count = 4096, Response Queue count = 512
Total number of active commands = 0
Total number of interrupts = 206796
    Device queue depth = 0x20
Number of free request entries = 2212
Number of mailbox timeouts = 0
Number of ISP aborts = 0
Number of loop resyncs = 0
Number of retries for empty slots = 0
Number of reqs in pending_q= 0, retry_q= 0, done_q= 0, scsi_retry_q= 0
Host adapter:loop state = <READY>, flags = 0x1e13
Dpc flags = 0x4080000
MBX flags = 0x0
Link down Timeout = 030
Port down retry = 035
Login retry count = 035
Commands retried with dropped frame(s) = 0
Product ID = 0000 0000 0000 0000


SCSI Device Information:
scsi-qla0-adapter-node=200000e08b854260;
scsi-qla0-adapter-port=210000e08b854260;
scsi-qla0-target-0=203500a0b81d2b95;
scsi-qla0-target-1=203400a0b81d2b95;
FC Port Information:
scsi-qla0-port-0=200400a0b81d2b95:203400a0b81d2b95:0000e4:0;
```

3. Check the output to verify connectivity with the controllers. The output in the previous example shows that one DS3000 with two targets **scsi-qla0-target-0** and **scsi-qla0-target-1**, with WWPN of **203500a0b81d2b95** and **203400a0b81d2b95**, are attached to the HBA.

Viewing the LUNs

View the LUNs by typing the following command:

```
# cat/proc/scsi/scsi
```

The LUN properties are displayed, as shown in the following example.

```
[root@x3650 /]# cat /proc/scsi/scsi
Attached devices:
Host: scsi2 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: 1726-4xx  FAStT  Rev: 0670
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi2 Channel: 00 Id: 01 Lun: 01
  Vendor: IBM      Model: 1726-4xx  FAStT  Rev: 0670
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi2 Channel: 00 Id: 01 Lun: 02
  Vendor: IBM      Model: 1726-4xx  FAStT  Rev: 0670
  Type:   Direct-Access                    ANSI SCSI revision: 05
.....
Host: scsi3 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: 1726-4xx  FAStT  Rev: 0670
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi3 Channel: 00 Id: 01 Lun: 01
  Vendor: IBM      Model: 1726-4xx  FAStT  Rev: 0670
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi3 Channel: 00 Id: 01 Lun: 02
  Vendor: IBM      Model: 1726-4xx  FAStT  Rev: 0670
  Type:   Direct-Access                    ANSI SCSI revision: 05
.....
Host: scsi4 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: VirtualDisk      Rev: 0670
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: VirtualDisk      Rev: 0670
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 02
  Vendor: IBM      Model: VirtualDisk      Rev: 0670
  Type:   Direct-Access                    ANSI SCSI revision: 05
```

The previous example shows three LUNs, which are assigned to a partition with two HBAs for dual-path, with RDAC installed.

Notes:

1. Each LUN is displayed on both of its configured paths. The active, assigned HBAs are numbers 2 and 3.
2. The **VirtualDisk** is the RDAC logical disk. There is one VirtualDisk for each LUN, regardless of the number of paths available.
3. If RDAC is installed, but you do not see any VirtualDisks in the output, the most likely cause is that the mpp_Vhba module is not 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 restart the host.

SLES 9 and SLES 10: On SLES operating systems only, you can use the **lsscsi** command to view LUNs, instead of **cat/proc/scsi/scsi**. The advantage of using the **lsscsi** command is that the output is easier to read, as shown in the following example.

```
[root@x3650 /]# lsscsi
[02:0:0:0]   disk    IBM    1726-4xx  FAStT   0270 -
[02:0:0:1]   disk    IBM    1726-4xx  FAStT   0270 -
[12:0:0:2]   disk    IBM    1726-4xx  FAStT   0270 -
...
[03:0:0:0]   disk    IBM    1726-4xx  FAStT   0270 -
[03:0:0:1]   disk    IBM    1726-4xx  FAStT   0270 -
[03:0:0:2]   disk    IBM    1726-4xx  FAStT   0270 -
...
[04:0:0:0]   disk    IBM    VirtualDisk       0270 /dev/sdb
[04:0:0:1]   disk    IBM    VirtualDisk       0270 /dev/sdc
[04:0:0:2]   disk    IBM    VirtualDisk       0270 /dev/sdd
```

The fourth column shows the DS3000 model type and the sixth column shows the Linux-assigned device number. In this example, only the Virtual RDAC devices are assigned a system device name. If RDAC is not installed, system device names are displayed for all of the devices listed in the output.

Associating the system LUN devices with DS3000 LUNs

To associate the system LUN devices with DS3000 luns, type the following command:

```
# SMdevices
```

The device properties are displayed, as shown in the following example.

```
[root@x3650 /]# /opt/IBM_DS3000/util/SMdevices
IBM System Storage DS4000/FAStT Storage Manager Devices, Version 09.17.A5.01
Built Wed Mar 28 11:53:31 CST 2007
(C) Copyright International Business Machines Corporation, 2003-2007 Licensed Ma
terial - Program Property of IBM.  All rights reserved.

  /dev/sdb (/dev/sg1) [Storage Subsystem DS3K_SATA_Test, Logical Drive data-1,
LUN 0,
Logical Drive ID <600a0b80001d2b9500004d80469bbb18>, Preferred Path
(Controller-A): In Use]
  /dev/sdc (/dev/sg2) [Storage Subsystem DS3K_SATA_Test, Logical Drive data-2,
LUN 1,
Logical Drive ID <600a0b80001d2b9500004d82469bbb20>, Preferred Path
(Controller-A): In Use]
  /dev/sdd (/dev/sg3) [Storage Subsystem DS3K_SATA_Test, Logical Drive data-3,
LUN 2,
Logical Drive ID <600a0b80001d2b9500004d7e469bbaa0>, Preferred Path
(Controller-A): In Use]
```

The previous example shows the assigned system device numbers that are associated with the DS3000 storage subsystem name, and storage assigned LUN ID, as they are displayed on the DS3000 storage subsystem.

On SLES and RHEL operating systems, you can use the following steps instead of using **SMdevices** to display device properties.

Type the following command from the linuxrdac directory:

```
# lsvdev
```

The following is an example of the **lsvdev** command output:

```
[root@x3650 mpp]# ./lsvdev
        Array Name      Lun    sd device
        ------------------------------------
        DS3K_SATA_Test   0     -> /dev/sdb
        DS3K_SATA_Test   1     -> /dev/sdc
        DS3K_SATA_Test   2     -> /dev/sdd
        DS3K_SATA_Test   3     -> /dev/sde
        DS3K_SATA_Test   4     -> /dev/sdf
        DS3K_SATA_Test   5     -> /dev/sdg
```

Performing common maintenance tasks

This section describes some common system administration tasks. The procedures in this section apply to both SLES and RHEL installations. The exceptions are noted.

Checking LUN size

To check the size of a LUN, complete the following steps:

1. Type the following commands:

   ```
   #cd /sys/block/sdxx
   # cat size
   ```

   where *xx* is the device name.

   A number is displayed, as shown 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)
   ```

   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

This section explains how to use dynamic volume expansion (DVE) with Linux on POWER.

Dynamic volume expansion (DVE) increases the size of a logical drive. 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 increase the capacity of the array by adding drives.

Before you begin: Make sure that there is sufficient available free capacity within the array. Free capacity can be determined using the DS3000 Storage Manager software. In the Summary/Arrays & Logical Drives view of the Subsystem Management window, expand the desired array and scroll to the bottom of the logical drives list. Available free capacity is an entry in the logical drives list.

Performing a dynamic capacity expansion operation

To perform a dynamic capacity expansion operation, complete the following steps:

1. In the Modify/Add Free Capacity (Drives) view, select the array whose capacity is to be increased.
2. Select **Next**.

3. Select an available drive in the “Add capacity to array” window.
4. Select **Finish**.

   Note: The available free disks and their capacity are displayed in the window.

After the operation is started, a clock is displayed beside every logical drive in the array and remains until the operation is completed. The process might take several hours and must be allowed to complete before continuing to the next steps.

Note: The time required for the operation to be completed depends on the I/O activity, capacity being added, drive technology, and other factors.

Performing a dynamic volume expansion operation

Performing a dynamic volume expansion requires the functionality of the Storage Manager software SMcli interface or the Storage Manager software script editor.

The following example shows the command syntax that is required to run the dynamic volume expansion from the script editor window on a logical drive named Engineering_FEB. The capacity added is 1 GB.

```
set logicalDrive ["Engineering_FEB"] addCapacity=1GB;
```

For more information about expanding Logical VolumeManager (LVM) logical drives, see the man page for the **chvg** command.

To perform a DVE on the Linux host, complete the following steps:

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 Storage Manager software SMcli interface or the Storage Manager software script editor, enter the following command:

   ```
   set logicalDrive [logical_drive_name] addCapacity=xGB;
   ```

   where *logical_drive_name* is the name of the logical drive that will increase and *x* is the amount to increase the logical drive.

   Note: You must wait for the process to complete before you can begin any host intervention. If the storage subsystem is busy, this process might take several hours to complete.
4. When the LVE process is complete, on the host, rescan the logical drive by typing the following commands:

   ```
   # cd /sys/block/sdxx/device
   # echo 1 > rescan

   where xx is the device name.
   ```
5. Check the size of the logical drive using the steps that are described in “Checking LUN size” on page 65.
6. Remount the logical drive.

Using the SMdevices utility to find information about LUNs

SMutil features an important utility, **SMdevices**, which you can use to map the host-assigned device name for each LUN back to its corresponding DS3000 storage subsystem device.

In the **SMdevices** output, you can view the DS3000 storage subsystem information, as it is shown on SMclient.

Note: The examples in the list refer to the sample **SMdevices** output.

The following example shows a sample **SMdevices** output for the subsystem **DS3K_SATA_Test**.

```
[root@x3650 /]# /opt/IBM_DS3000/util/SMdevices
IBM System Storage DS4000/FAStT Storage Manager Devices, Version 09.17.A5.01
Built Wed Mar 28 11:53:31 CST 2007
(C) Copyright International Business Machines Corporation, 2003-2007 Licensed Ma
terial - Program Property of IBM.  All rights reserved.

  /dev/sdb (/dev/sg1) [Storage Subsystem DS3K_SATA_Test, Logical Drive data-1, LUN 0,
Logical Drive ID <600a0b80001d2b9500004d80469bbb18>, Preferred Path (Controller-A):
In Use]
  /dev/sdc (/dev/sg2) [Storage Subsystem DS3K_SATA_Test, Logical Drive data-2, LUN 1,
Logical Drive ID <600a0b80001d2b9500004d82469bbb20>, Preferred Path (Controller-A):
In Use]
  /dev/sdd (/dev/sg3) [Storage Subsystem DS3K_SATA_Test, Logical Drive data-3, LUN 2,
Logical Drive ID <600a0b80001d2b9500004d7e469bbaa0>, Preferred Path (Controller-A):
In Use]
```

- Host assigned name: `/dev/sdb`
- DS3000 storage subsystem name: `DS3K_SATA_Test`
- Logical drive name: `data-1`
- LUN ID: `LUN 0`
- Preferred controller owner, and whether that controller is currently controlling the logical drive

Viewing or changing queue depth and the no device timeout

Setting the *queue_depth* attribute to the applicable value is important for system performance. If you have a large DS3400 configuration with many logical drives and hosts attached, use this setting for high performance.

Set the “no device timeout” option to 144 seconds for System p servers that are running a Linux operating system and are using Emulex FC HBAs. Changing this parameter can improve error handling when RDAC is used. See the Emulex readme file for information about preferred settings for non-RDAC configurations.

This section provides methods for calculating your system 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

On DS3400 storage subsystems, use the following formula to calculate the maximum queue depth for your system:

2048 / (number-of-hosts × LUNs-per-host)

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

2048 / (4 × 32) = 16

Important:

- 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.

- 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.

Viewing the current queue depth

To see the queue depth of each LUN, type the following command:

```
# cat /proc/scsi/sg/{device_hdr,devices}
```

Viewing the current no device timeout

To see the "no device timeout" of each HBA, type the following command:

```
# cat /sys/class/scsi_host/hostx/lpfc_nodev_tmo
```

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
```

Changing the queue depth and no device timeout

To change the queue depth and the "no device timeout", complete the following steps:

1. Using the vi editor, open one of the following configuration files for editing, according to your operating system:
   - **SLES 9 and SLES 10:**
     /etc/modprobe.conf.local
   - **RHEL 4 and RHEL 5:**
     /etc/modprobe.conf
2. Add one of the following option parameters to the configuration file, according to the type of HBA:
   - **Emulex-based HBA:** options lpfc_lun_queue_depth=*xx* lpfc_nodev_tmo=144
   - **QLogic-based HBA:** ql2xmaxqdepth=*xx*

   where *xx* is the new queue depth value.
3. Set the new queue depth value , as follows:
   - **With RDAC installed:**
     Type the **mppUpdate** command, then restart the server.
   - **Without RDAC installed:**
     Uninstall the applicable HBA driver, and then reload it.

Example: To change the queue depth value to 10 and the "no devices timeout" to 144 on an RHEL 4 system with an Emulex-based HBA and a QLogic-based HBA, type the following commands:

```
# vi /etc/modprobe.conf
options lpfc lpfc_lun_queue_depth=10 lpfc_nodev_tmo=144
options qla2xxx ql2xmaxqdepth=10
```

Replacing IBM host bus adapters on a system

This section provides requirements and procedures for replacing IBM host bus adapters in System p servers that are running SLES 9 SP4, SLES 10 SP1, RHEL 4.6, or RHEL 5.1 by using the PCI Hotplug tools.

Requirements

PCI Hotplug tools:

Make sure that the following tools are installed in the /usr/sbin directory:

- **lsslot**
- **drslot_chrp_pci**

To install these tools, complete the following steps:

1. Make sure 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 one of the following Web sites:

   SLES 9:

   https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/suselinux/other/

   SLES 10:

   https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/suselinux/other/sles10.html

   RHEL 4:

   https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/redhat/other/rhel4.html

   RHEL 5:

   https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/redhat/other/rhel5.html

   To navigate to these addresses, complete the following steps:

   a. Go to http://www.ibm.com/systems/support/.
   b. Under **Product support**, click **Power**.
   c. Under **Popular links**, click **Firmware updates**.
   d. Under **Other microcode-related services**, click **Download hardware service aids for Linux on Power**.
   e. Under **Red Hat** or **Suse Linux**, click **on other servers**.
   f. Click the tabs to display the lists of applicable service and productivity tools.
3. Download the rpm files for the following tools:
   - Platform Enablement Library
   - PCI Hotplug Tools
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:

Host bus adapters that are installed in blade servers are not supported by this procedure.

PCI core:

To make sure that the PCI core is loaded on the system, type the following command:

```
# ls -l /sys/bus/pci/slots
```

If the PCI core is loaded, the output looks similar to the following example.

```
elm17c224:/usr/sbin # ls -l /sys/bus/pci/slots
total 0
drwxr-xr-x  8 root root 0 Sep  6 04:29 .
drwxr-xr-x  5 root root 0 Sep  6 04:29 ..
drwxr-xr-x  2 root root 0 Sep  6 04:29 0000:00:02.0
drwxr-xr-x  2 root root 0 Sep  6 04:29 0000:00:02.4
drwxr-xr-x  2 root root 0 Sep  6 04:29 0000:00:02.6
drwxr-xr-x  2 root root 0 Sep  6 04:29 0001:00:02.0
drwxr-xr-x  2 root root 0 Sep  6 04:29 0001:00:02.6
drwxr-xr-x  2 root root 0 Sep  6 04:29 control
```

If the /sys/bus/pci/slots directory is not displayed, the PCI core is not loaded.

rpaphp driver:

To make sure that the rpaphp driver is loaded on the system, type the following command:

```
ls -l /sys/bus/pci/slots/*
```

If the rpaphp driver is loaded, the output looks similar to the following example.

```
elm17c224:/usr/sbin # ls -l /sys/bus/pci/slots/*
/sys/bus/pci/slots/0000:00:02.0:
total 0
drwxr-xr-x  2 root root    0 Sep  6 04:29 .
drwxr-xr-x  8 root root    0 Sep  6 04:29 ..
-r--r--r--  1 root root 4096 Sep  6 04:29 adapter
-rw-r--r--  1 root root 4096 Sep  6 04:29 attention
-r--r--r--  1 root root 4096 Sep  6 04:29 max_bus_speed
-r--r--r--  1 root root 4096 Sep  6 04:29 phy_location
-rw-r--r--  1 root root 4096 Sep  6 04:29 power
```

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 the **lsslot** tool, and provides examples.

Syntax for the lsslot command: The **lsslot** tool has the following syntax:

- 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 that 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 *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 -a
```

The resulting output looks similar to the following example.

```
# 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 empty PCI Hotplug slots: Type the following command to list all empty PCI Hotplug slots:

```
# lsslot -c pci -a
```

The resulting output looks similar to the following example.

```
# 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
```

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 example.

```
# 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
```

Show detailed information about a particular device: Select a device number from the output of the **# lsslot -c pci -o** command, as shown in the previous output example. Type the following command to show detailed information about that 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 example.

```
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: The **drslot_chrp_pci** tool has the following syntax:

- 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: To replace a PCI Hotplug HBA by using the drslot_chrp_pci command, complete the following steps.

Note: In the following procedures, the variable *slot-name* is the slot that contains the HBA that you are replacing.

Attention: Before you remove the HBA, you must remove the Fibre Channel cable that is attached to it. The Fibre Channel cable must remain unattached for at least 5 minutes to make sure that all I/O activity is transferred to the alternate path. Failure to remove the Fibre Channel cable might have unintended results.

1. To identify the PCI Hotplug slot, type the following command.

   ```
   # drslot_chrp_pci -i -s  slot-name
   ```

   where *slot-name* is the name of the slot for the HBA you are replacing. (For example, U7879.001.DQD014E-P1-C3.)

   The LED at slot *slot-name* begins flashing and the following message is displayed:

```
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. To hot unplug the HBA from the slot, complete the following steps:
   a. Remove the Fibre Channel cable that is connected to this HBA and wait for failover to be completed.
   b. After failover is completed, type the following command:

      ```
      # drslot_chrp_pci -r -s   slot-name
      ```

      The following message is displayed:

```
The visual indicator for the specified
PCI slot has been set to the identify
state. Press Enter to continue or
enter x to exit.
```

   c. Press Enter. The following message is displayed:

```
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.
```

   d. Press Enter.

e. Physically remove the HBA from the slot.

f. Type the following command to make sure that the slot is empty:

```
# lsslot -c pci -s    slot-name
```

If the slot is empty, the resulting output looks similar to the following example.

```
# Slot                  Description                         Device(s)
U7879.001.DQD014E-P1-C3  PCI-X capable, 64 bit, 133MHz slot  Empty
```

3. To hot plug the HBA into the slot, complete the following steps:

a. Type the following command:

```
# drslot_chrp_pci -a -s      slot-name
```

The following message is displayed:

```
The visual indicator for the specified
PCI slot has been set to the identify
state. Press Enter to continue or
enter x to exit.
```

b. Press Enter. The following message is displayed:

```
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.
```

c. Insert the new HBA into the slot.

d. Type the following command to make sure that the slot is not empty:

```
# lsslot -c pci -s     slot-name
```

If the slot is not empty, the resulting output looks similar to the following example.

```
# 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 DS3000 storage subsystem

After you replace the HBA, to reassign its new WWPN on the DS3000 storage subsystem, complete the following steps:

1. If RDAC is installed, type the following command to recognize the new HBA:

```
# mppBusRescan
```

2. Start the Storage Manager software.
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. Make sure that the newly-replaced HBA recognizes the LUNs by checking the **lsscsi** command output.

Chapter 5. Completing the Storage Manager software installation and configuration

This chapter contains information about using the Enterprise Management and Subsystem Management features of the Storage Manager 2 software to complete the installation tasks for the AIX and Linux on POWER operating system environments.

To complete the Storage Manager 2 software installation, perform the following tasks. Each of these tasks is discussed in detail in the following sections.

1. Perform a discovery of storage subsystems.
2. Add devices, if necessary.
3. Set up alert notifications, if you want to.
4. Name the storage subsystems.
5. Download firmware and NVSRAM.
6. Create arrays and logical drives.
7. Configure the heterogeneous hosts.
8. Map logical drives to a partition.
9. Perform other storage subsystem tasks.

Starting the Storage Manager software and discovering storage subsystems

To start the Storage Manager software and discover storage subsystems, complete the following steps:

1. Choose one of the following commands:
   - **For AIX:** Enter `/usr/SMsmclient/SMclient`
   - **For Linux:** Enter `/opt/IBM_DS3000/client/SMclient` (by default).

   The client software starts and the Enterprise Management window opens. Then, the Task Assistant window (shown in Figure 8 on page 79) and the Select Addition Method window (shown in Figure 6 on page 76) open on top of the Enterprise Management window.

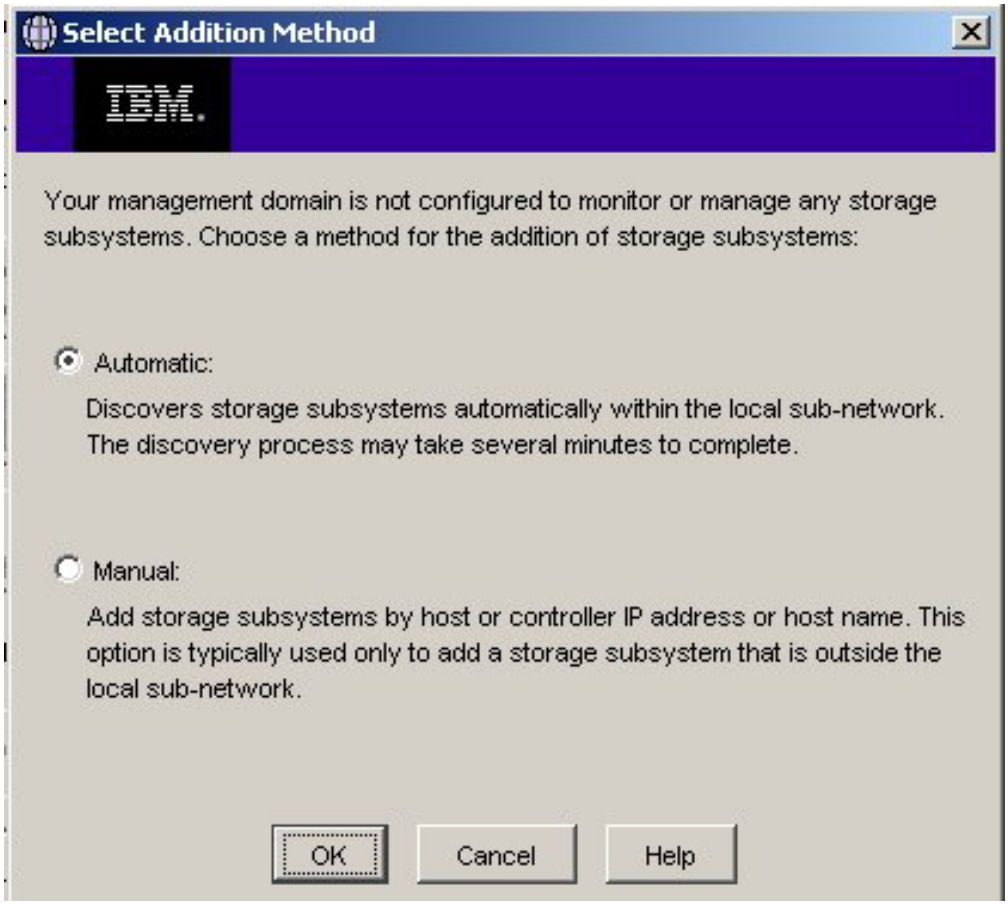

Figure 6. Select Addition Method window

2. Select **Automatic** to discover all hosts and storage subsystems that are attached to the local subnetwork. Select **Manual** to specify the IP address of a device that is attached to the local subnetwork.

   Note: The Enterprise Management window can take up to 1 minute to refresh after an initial automatic discovery.

3. Make sure that each host and storage subsystem is displayed in the Enterprise Manager window.

   If a host or storage subsystem is not displayed, complete the following tasks:

   - Check the hardware and hardware connections for possible problems (see the hardware documentation for specific procedures).
   - For more information about discovering storage subsystems, see the Enterprise Management online help.
   - If you are using the direct-management method, make sure that all hosts and storage subsystems are connected to the same subnet network. If you are using the host-agent method, make sure that the I/O connection between the host and the storage subsystem is made.
   - Make sure that all of the preparation steps for setting up the storage subsystem for a direct-managed system in Chapter 2, “Preparing for installation,” on page 15 are completed. If they are, use the Add Device option to add the IP addresses of the storage subsystem. Add both IP addresses of the controllers. Otherwise, a `partially-managed device` error message is displayed when you try to manage the storage subsystem.

- If you are using the host-agent management method, complete the following steps:
  a. Make sure that the SMagent is installed on the host.
  b. Make sure that you have a valid I/O connection from the storage subsystems to the host in which the SMagent is installed.
  c. Make sure that all of the preparation steps that are outlined in Chapter 2, “Preparing for installation,” on page 15 are complete. If they are, complete the following steps:
     1) Run the hot_add utility.
     2) Restart the SMagent, using the information in “Stopping and restarting the host-agent software” on page 107.
     3) On the Enterprise Management window, click the host and click **Tools → Rescan**.

Note: In certain situations, a storage subsystem might be duplicated in the device tree after an automatic discovery. You can remove a duplicate storage management icon from the device tree by using the Remove Device option in the Enterprise Management window.

4. Make sure that the status of each storage subsystem is Optimal. If a device shows a status of Unresponsive, right-click the device and select **Remove Device** to delete it from the management domain. Then, use the Add Device option to add it to the management domain again. For instructions for removing and adding devices, see the Enterprise Management window online help.

   Figure 7 shows the Enterprise Management window after an initial automatic discovery.

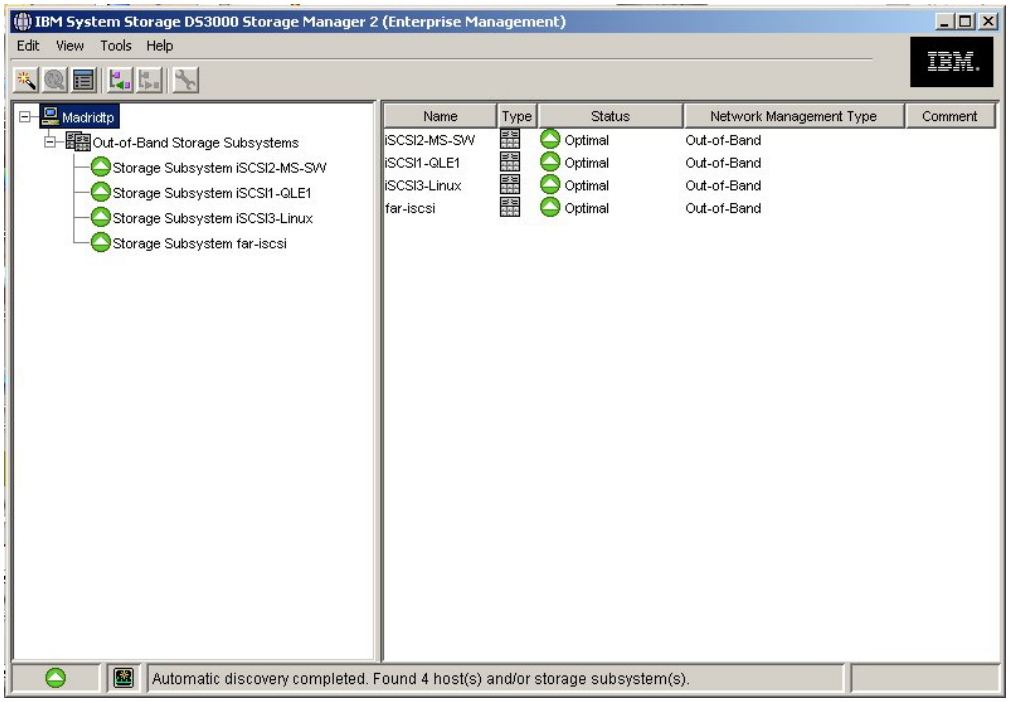

Figure 7. Enterprise Management window

Using the Task Assistant

When you start the Storage Manager software, the client software starts, and the Enterprise Management window opens. Then, the Task Assistant window opens and provides a central location from which you can choose to perform the most common tasks.

Note: The Task Assistant window opens automatically each time you open the Enterprise Management window, unless you select the **Don't show the task assistant at start-up again** check box at the bottom of the window.

The Task Assistant provides shortcuts to the following tasks:

- Adding storage subsystems
- Naming or renaming storage subsystems
- Configuring alerts
- Managing storage subsystems

Figure 8 on page 79 shows the Enterprise Management Window Task Assistant.

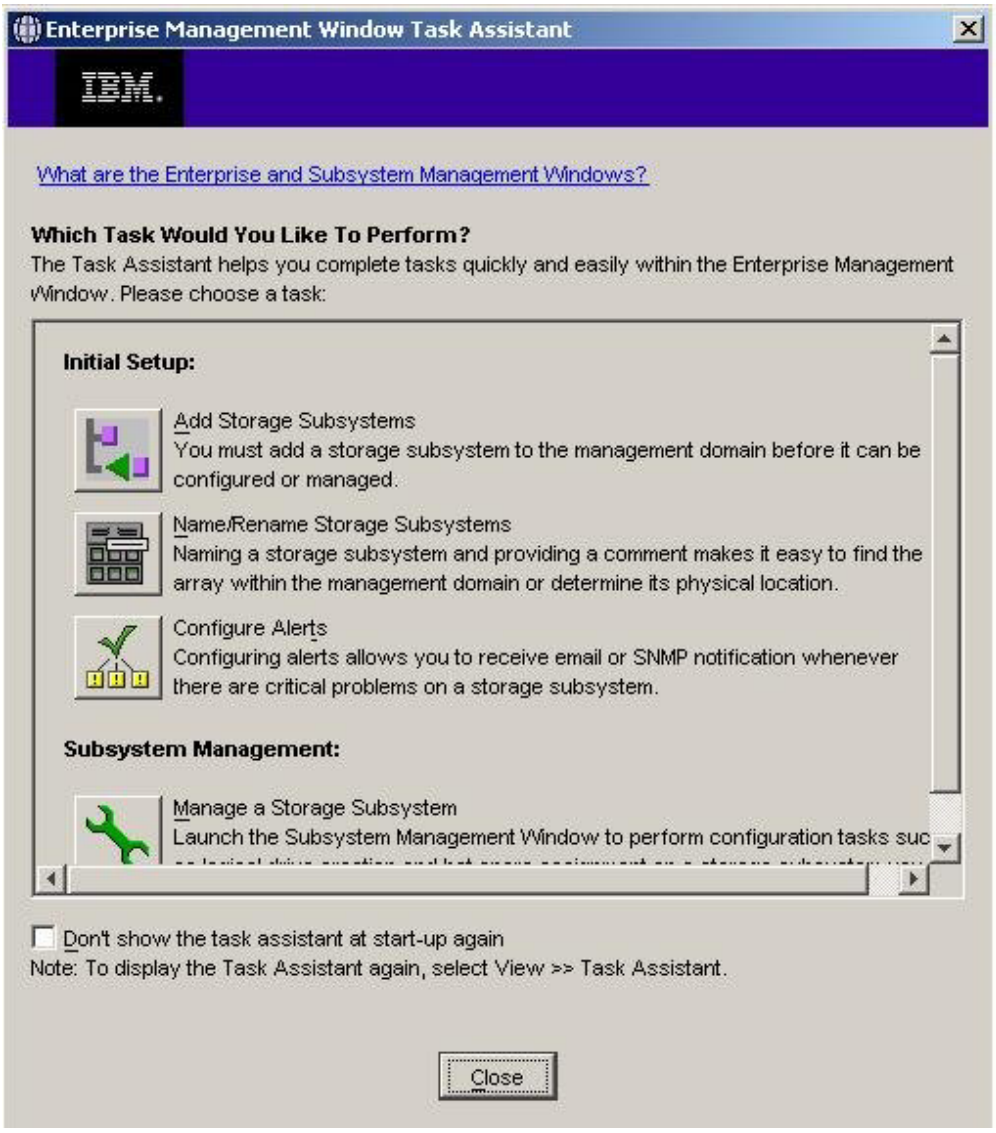

Figure 8. Enterprise Management Window Task Assistant

To open the Task Assistant after you have started the Storage Manager software, click **View → Task Assistant** from the Enterprise Management window, or on the toolbar, click the **Task Assistant** icon  .

Adding storage subsystems

To manually add more hosts or storage subsystems, from the Enterprise Management window, click **Edit → Add Storage Subsystem**. You can use this option to selectively manage a group of storage subsystems from a Storage Manager Client, and you can add devices that are to be managed and were not discovered during the initial discovery. For more information, see the Enterprise Management window online help.

Important:

- When you add an unidentified device, a window opens confirming that the device was successfully added and that you can manage the storage subsystem. Any

device that is listed in the unidentified device group will not be able to be managed until it is discovered correctly. Make sure that the connections are made, and try to add the device again.

- When you add new storage subsystems to the existing storage subsystems in a SAN that are managed through the host-agent software, you must stop and restart the host-agent service. When the host-agent service restarts, the new storage subsystem is detected. For more information, see “Stopping and restarting the host-agent software” on page 107. Then, go to the Enterprise Management window and click **Tools → Rescan** to add the new storage subsystems to the management domain.
- When you add new storage subsystems to existing storage subsystems that are managed through the direct (out-of-band) management method, be sure to specify the IP addresses for both controllers.

Naming storage subsystems

To name your storage subsystem, complete the following steps:

1. In the Enterprise Management window, select a storage subsystem, right-click the storage subsystem and select **Rename**. The Rename Storage Subsystem window opens.

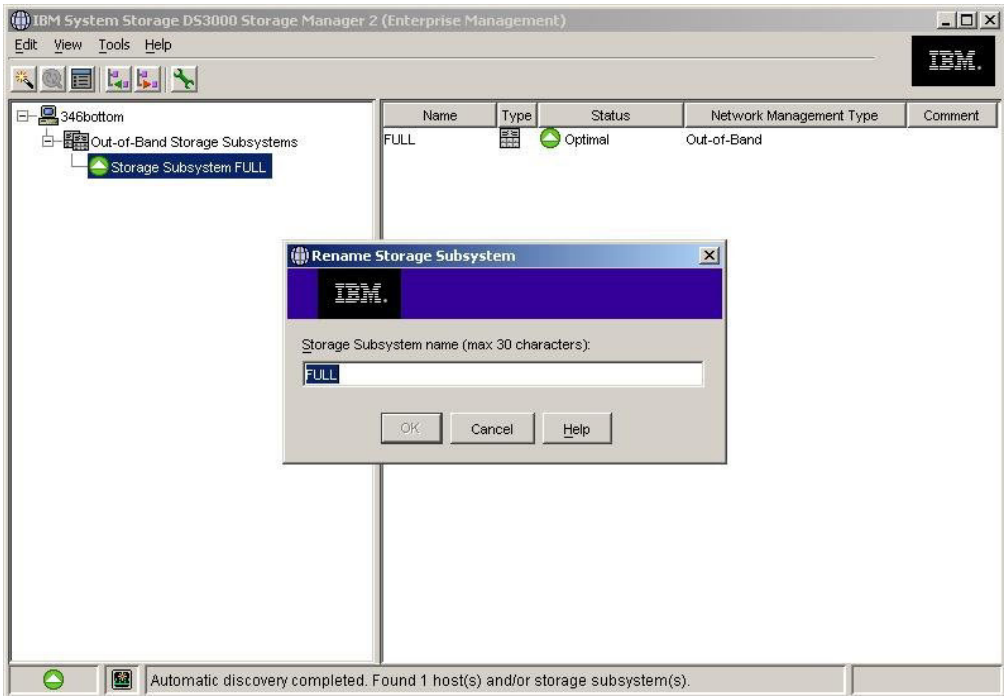

Figure 9. Rename Storage Subsystem window

2. Type the name of the storage subsystem (see Table 9 on page 113). Then, click **OK**. A warning message is displayed stating that changing the storage subsystem name can cause host applications to lose access to the storage subsystem. Click **Yes** to continue or **No** to cancel naming the storage subsystem.
3. Repeat this procedure for each unnamed storage subsystem. For more information, see the topic about renaming storage subsystems in the Subsystem Management window online help.

Configuring alerts

After you add devices to the management domain, you can set up alert notifications to report critical events on the storage subsystems. The following alert-notification options are available:

- Notification to a designated network management station (NMS) through Simple Network Management Protocol (SNMP) traps
- Notification to designated e-mail addresses

Note: You can monitor storage subsystems only within the management domain. If you do not install the Event Monitor service, the Enterprise Management window must remain open. If you close the window, you will not receive any alert notifications from the managed storage subsystems. For more information, see the Enterprise Management window online help.

To set up alert notification, from the Enterprise Management window, click **Edit → Configure Alerts**.

To set up alert notification to a network management station (NMS) by using SNMP traps, complete the following steps:

1. Insert the host kit support CD into the CD drive on an NMS. You must set up the designated management station only once.
2. Copy the SM2.MIB file from the SM2MIB directory to the NMS.
3. Follow the steps that are required by your NMS to compile the management information base (MIB) file. For more information, contact your network administrator or see the documentation that comes with your storage-management software.

The Initial Setup Tasks window

From the Subsystem Management window, you can open the Initial Setup Tasks window, which contains links to the following tasks.

- Locating the storage subsystem
- Renaming the storage subsystem
- Setting the storage subsystem password
- Configuring host access
- Managing iSCSI settings (DS3300 only)
- Configuring the storage subsystem

To open the Initial Setup Tasks window, complete the following steps:

1. In the Enterprise Management window, double-click the storage subsystem for which you want to perform the setup tasks. The Subsystem Management window opens.
2. Click the **Summary** tab.
3. On the **Summary** page, click **Perform Initial Setup Tasks**. The Initial Setup Tasks window opens, as shown in Figure 10 on page 82.

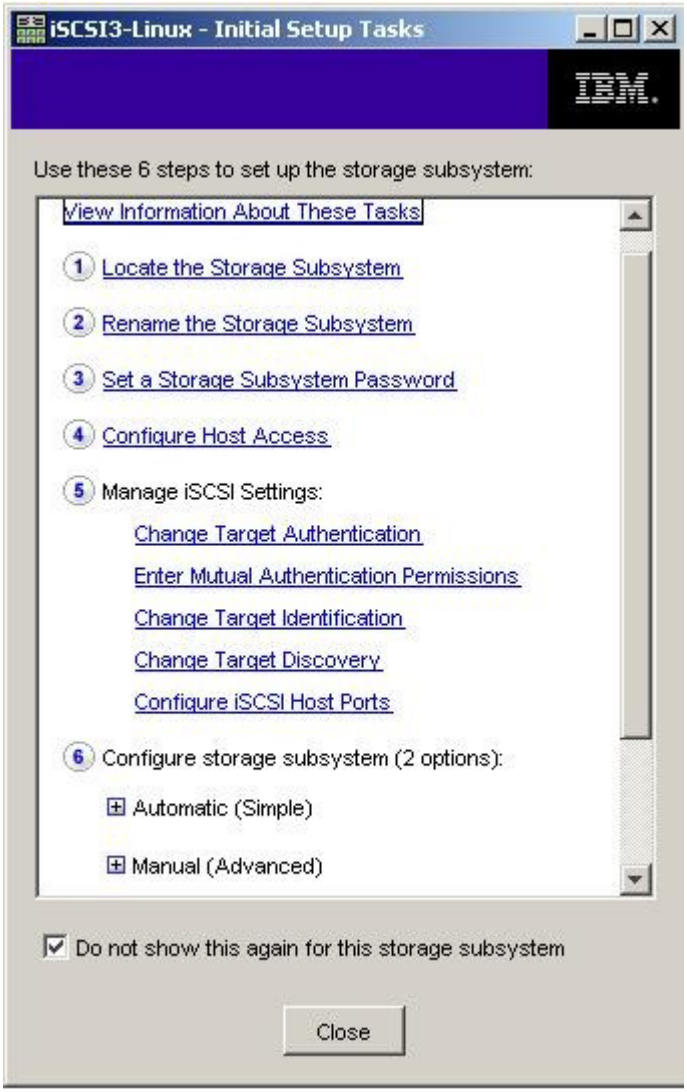

Figure 10. Initial Setup Tasks window in the Subsystem Management window

Note: The Initial Setup Tasks window opens automatically each time you open the Subsystem Management window, unless you select the **Do not show this again for this storage subsystem** check box at the bottom of the window.

Creating a storage subsystem profile

Important: Create a storage subsystem profile and save it in a safe place whenever you modify the arrays and logical drives in the storage subsystem. This profile contains detailed controller information, including logical and physical disk configuration information, that you can use to help recover the configuration in the event of a failure. Do not save the profile in the logical drives that are created in the DS3000 storage subsystem whose profile was collected and saved.

To save the storage subsystem profile, complete the following steps:

1. On the Subsystem Management window, click **Summary → Storage Subsystem Profile**. The Storage Subsystem Profile window opens.
2. Click **Save As**.

3. Select the sections that you want to save, and provide the file name in which to save the Profile.

You can also click **Support → Gather Support Information**, as shown in Figure 11, to collect the inventory, status, diagnostic, and performance data from the storage subsystem and save them to a single compressed file.

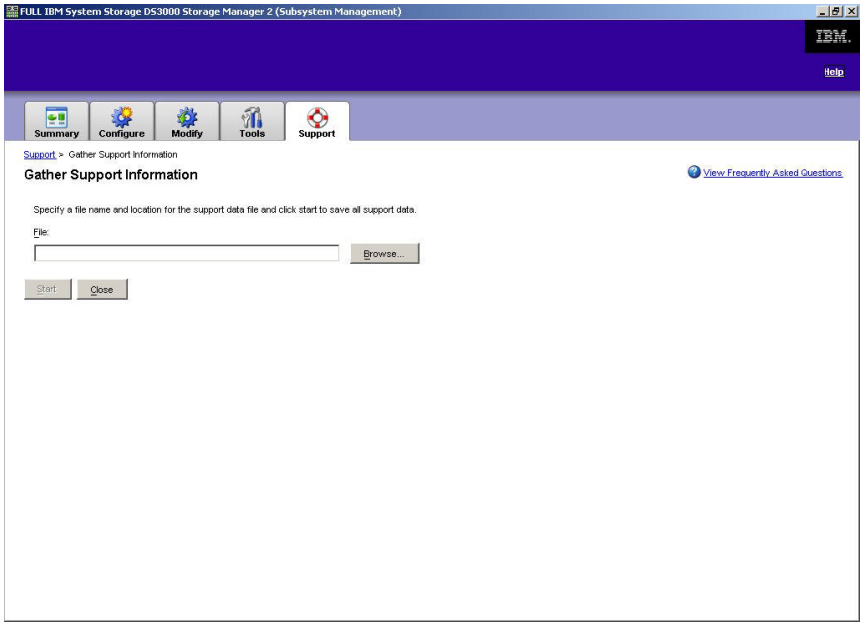

Figure 11. Gather Support Information window

Configuring host access

Before you use the logical drives in a host server, you must define and configure a host for access.

To configure host access, complete the following steps:

1. Click the **Configure** tab in the Subsystem Management window. Choose one of the following options:
   - To configure host access automatically, click **Configure Host Access (Automatic)** and follow the wizard instructions. After you have completed the instructions in the wizard, you are finished with this procedure.

     Note: To detect hosts automatically, the Storage Manager 2 Agent service must be running. Make sure that SMagent is installed before you proceed.
   - To configure host access manually, click **Configure Host Access (Manual)** and continue to step 2.
2. When the Configure Host Access (Manual) wizard starts, specify a host name and host type, and then click **Next**.

   Note: Before you use the logical drives in a host server, you must specify the correct host type. The host type determines how the storage subsystem controllers will work with each operating system on the hosts to which the controllers are connected. For the supported host types, see the Storage Manager software readme file.

3. When the **Specify HBA Host Ports** area is displayed, add the known or new host bus adapters in the **Selected HBA host port identifiers/aliases** field, and then click **Next**.
4. (Optional) When the **Specify Host Group** area is displayed, specify a host group that will share access with logical drives. When you specify a host group, you can create a new host group or select an existing host group that has already been defined. Select **No** or **Yes**, and then click **Next**.

Defining host groups

A *host group* is an entity in the Storage Partitioning topology that defines a logical collection of host servers that require shared access to one or more logical drives. You can grant individual hosts in a defined host group access to additional logical drives that are not shared by other nodes, but this requires an additional storage partition. You can make logical drive-to-LUN mappings to the host group or to an individual host in a host group.

Important: If you have purchased a Storage Partitioning premium feature, make sure that the premium feature is enabled. See the instructions for enabling premium features that came with your key, or contact your technical-support representative if you are unable to enable a premium feature.

Note: You can configure heterogeneous hosts when you define a host group. This enables hosts running different operating systems to access a single storage subsystem.

To define a host group, complete the following steps:

1. On the Subsystem Management window, click **Configure → Create Host Group** to start the Create Host Group wizard.
2. When the **Create Host Group** area is displayed, enter a name for the host group and add any available hosts to the host group.
3. Click **OK**.

   Note: To add hosts to the host group at this time, available hosts must already be configured for access. You can first create a host group without adding hosts and then add hosts later, using the Create Host Access (Manual) wizard. For more information, see “Configuring host access” on page 83.

Creating arrays and logical drives

A redundant array of independent disks (RAID) *array* is a set of hard disk drives that are logically grouped together.

A *logical drive* is a logical structure that is the basic structure that you create to store data on the storage subsystem. The operating system recognizes a logical drive as a single drive. Choose a RAID level to meet application needs for data availability and to maximize I/O performance.

The maximum number of drives that are supported in an array is 30. Each array can be divided into 1 through 256 logical drives.

To create arrays and logical drives, complete the following steps:

1. From The Enterprise Management window, double-click a storage subsystem to open a Subsystem Management window.
2. Click the **Configure** tab.
3. Click **Create Logical Drives**. The Create Logical Drives wizard starts and the Select Capacity Type window opens.

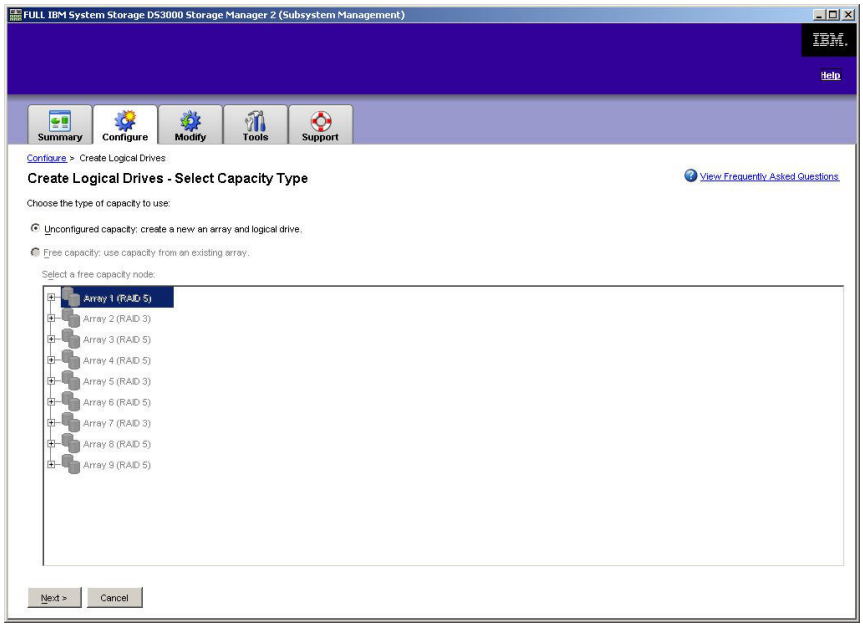

Figure 12. Create Logical Drives wizard window

4. Select **Unconfigured capacity** or **Free capacity**, and then click **Next**.

   Note: When you select **Free capacity**, you must first select the free capacity node from an existing array before you proceed with the wizard.

   If you select **Unconfigured capacity**, skip to step 10 on page 87.
5. From the **Select drive type** list, select **SAS** or **SATA**.

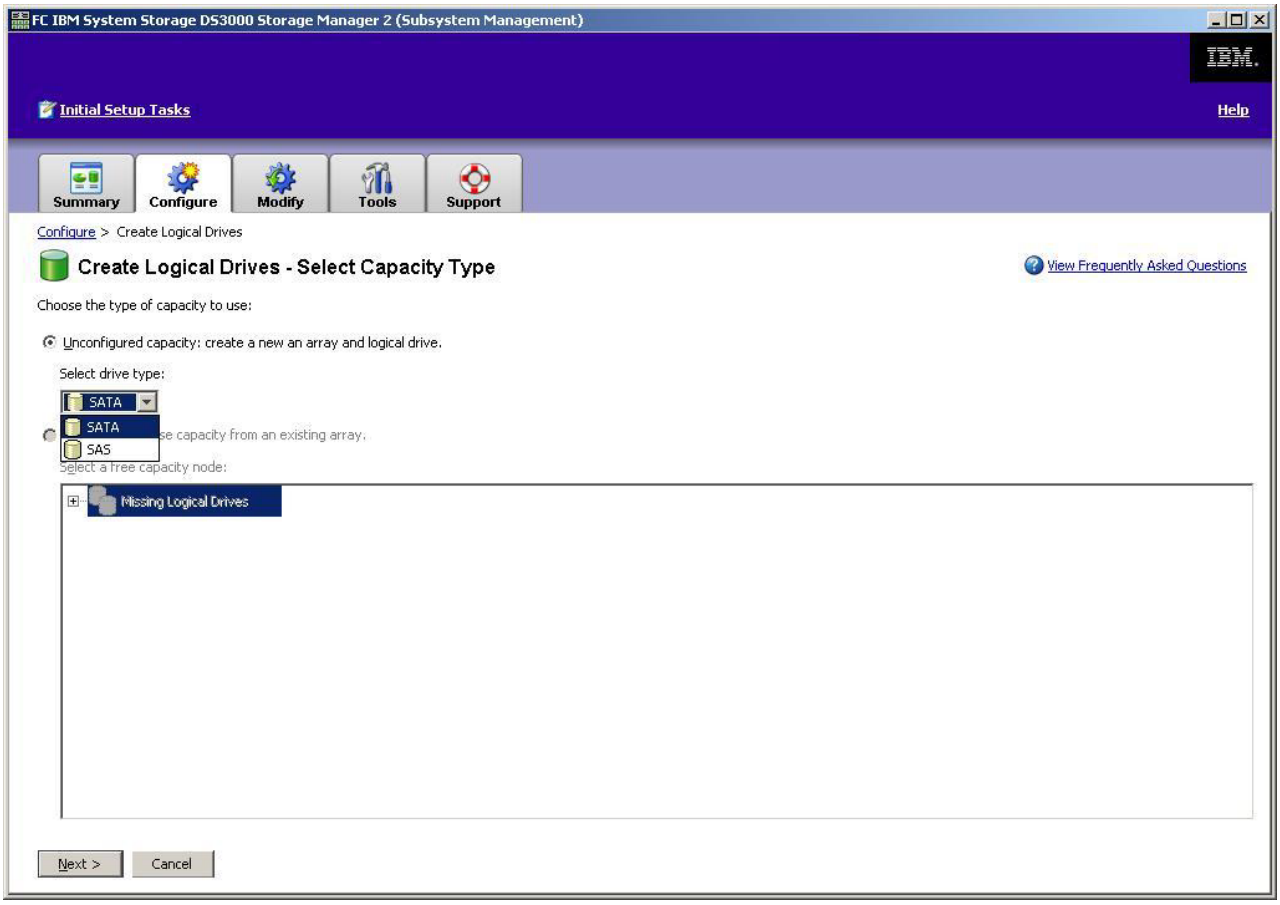

Figure 13. Selecting the capacity type

Note: You must create arrays by using drives that have the same disk type. You cannot mix both SATA and SAS drives within the same array.

6. When the **Drive Selection Choices** area is displayed, select either **Automatic** or **Manual** and click **Next**. If you select **Automatic**, continue to step 7. If you select **Manual**, skip to step 8.

   Note: Selecting **Automatic** is preferred for quick and easy configuration. This method enables you to choose from a list of automatically generated drive and capacity options. Selecting **Manual** gives advanced users a method for specifying drives to create a new array.

7. When the **Choose Configuration (Automatic Configuration)** area is displayed, select a RAID level and click **Finish**.

   Note: Depending on the selected RAID level, the automatic configuration configures all remaining unconfigured capacity that is currently available on the storage subsystem. Check the **Configuration summary** field for details about what will be configured.

8. When the **Manual Drive Selection** area is displayed, as shown in Figure 14 on page 87, select the RAID level. Select the drives from the **Unselected drives** area and click **Add** to move the drives to the **Selected drives** area. In the **Selected drives** area, click **Calculate Capacity**, and then click **Next**.

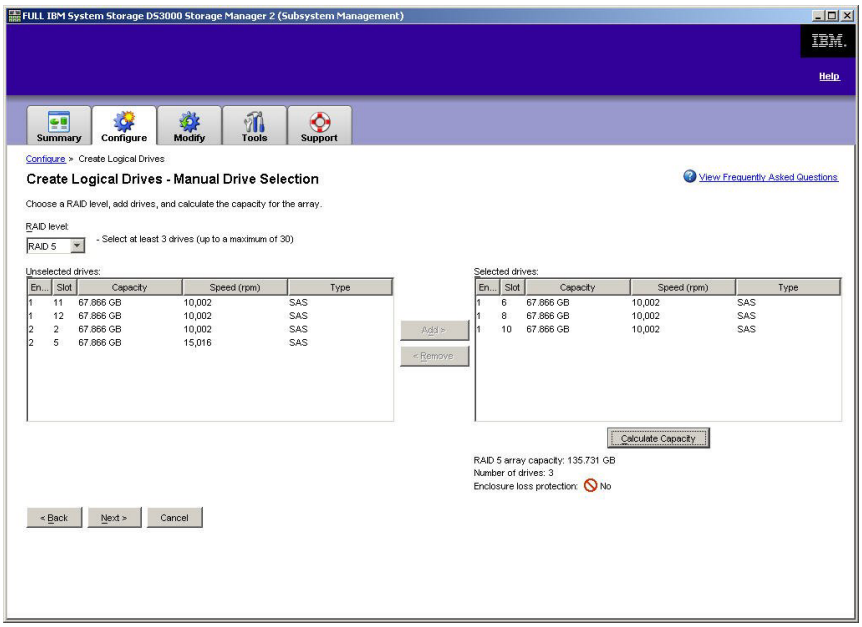

Figure 14. Manual drive selection area

9. When the **Specify Logical Drive** area is displayed, specify the capacity, name, and I/O characteristics of the logical drive, and click **Next**.
10. When the **Map Logical Drive To Host** area is displayed, select **Map now** to map a logical drive immediately, using the host that you defined in "Configuring host access" on page 83 or **Map later** to map a logical drive at a later time, and then click **Next**.

    Notes:

    a. If you did not configure hosts previously, a warning message is displayed. Read the message and click **OK**; then, continue with the Create Logical Drives wizard.

    b. To map the logical drive immediately, you must have previously configured a host or host group (be sure to assign a LUN number for your logical drive before you continue). If you select **Map later**, you can map the logical drive through the Create Host-to-Logical Drive Mappings wizard. For more information, see "Mapping LUNs to a partition."

Note: For cluster configurations, if you add or delete logical drives, you must make them known to both nodes A and B.

Mapping LUNs to a partition

Logical unit numbers (LUNs) are assigned to each logical drive when they are mapped to a partition.

To map LUNs to a new or existing partition, complete the following steps:

1. On the Subsystem Management window, click **Configure → Create Host-to-Logical Drive Mappings** to start the wizard.
2. When the **Select Host** area is displayed, select a host or host group and click **Next**.

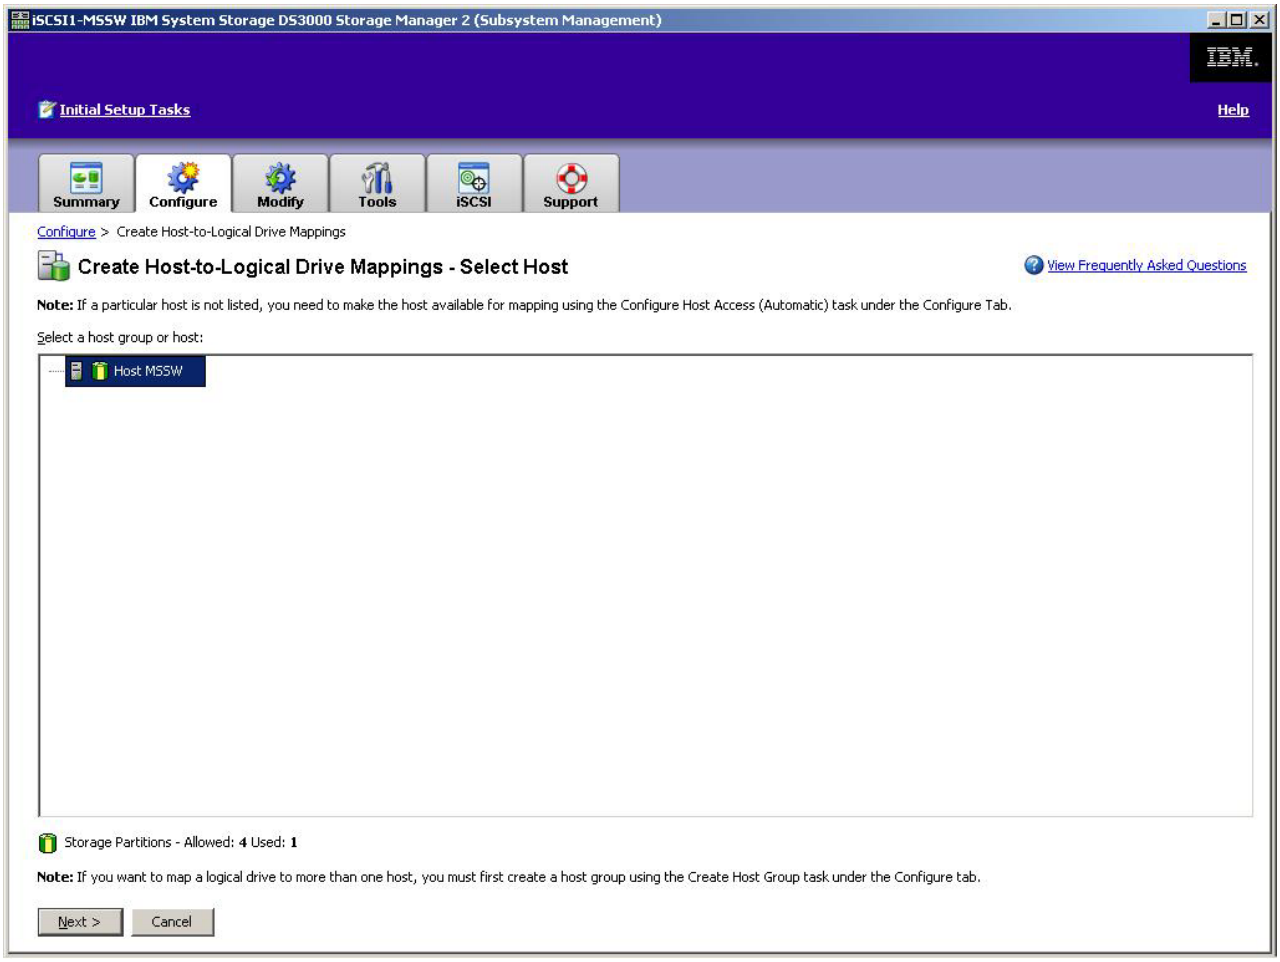

Figure 15. Select host area

3. When the **Select Logical Drives** area is displayed, select the logical drives that you want to map to the specified host, and click **Finish**.

Configuring hot-spare devices

You can assign available physical drives in the storage subsystem as hot-spare drives to keep data available. A hot spare is a drive that contains no data and that acts as a standby in case a drive fails in an array. If the drive in an array fails, the controllers automatically use a hot-spare drive to replace the failed drive while the storage subsystem is operating. The controller uses redundancy data to automatically reconstruct the data from the failed drive to the replacement (hot-spare) drive. This is called *reconstruction*.

The hot-spare drive adds another level of redundancy to the storage subsystem. If a drive fails in the storage subsystem, the hot-spare drive is automatically substituted without requiring a physical swap. If the hot-spare drive is available when a logical drive fails, the controller uses redundancy data to reconstruct the data from the failed logical drive to the hot-spare drive. When the failed drive is replaced or another drive is designated as a replacement, the data from the hot-spare drive is copied back to the replacement drive. This is called *copyback*.

To assign a hot-spare drive, open the Subsystem Management window, click the **Configure** tab, and select one of the following options:

- **Automatically assign drives**. If you select this option, hot spare drives are automatically created for the best hot spare coverage using the drives that are available.
- **Manually assign individual drives**. If you select this option, hot spare drives are created out of available drives that are manually selected from the Configure Hot Spares window.

If you choose to manually assign the hot-spare drives, select a drive with a capacity equal to or larger than the total capacity of the drive you want to cover with the hot spare. For maximum data protection, do not assign a drive as a hot-spare unless its capacity is equal to or greater than the capacity of the largest drive on the storage subsystem. Use only the largest capacity drives for hot-spare drives in mixed capacity hard disk drive configurations.

Replacing a hard disk drive

If a hard disk drive fails and a hot-spare drive is available, the hot spare begins a reconstruction process to temporarily replace the failed drive. With the configurable drive replacement policy, you can choose one of the following options:

- Replace the failed drive with a replacement drive. The replacement drive performs a copyback operation from the hot-spare drive and becomes part of the array.
- From the Replace Drive window (on the **Modify** tab), select the hot-spare drive that is currently part of the degraded array, to become a permanent member of the array.
- From the Replace Drive window (on the **Modify** tab), select another available drive in the storage enclosure to become designated as the replacement drive. This drive performs a copyback operation from the hot spare and becomes part of the array.

Note: The selection of a replacement drive cannot be completed until a hot-spare drive has completed its reconstruction process.

Managing iSCSI settings (DS3300 only)

Note: The **iSCSI** tab on the Subsystem Management window is available only for DS3300 storage subsystems.

On the Subsystem Management window, click the **iSCSI** tab. A window similar to the one in the following illustration opens.

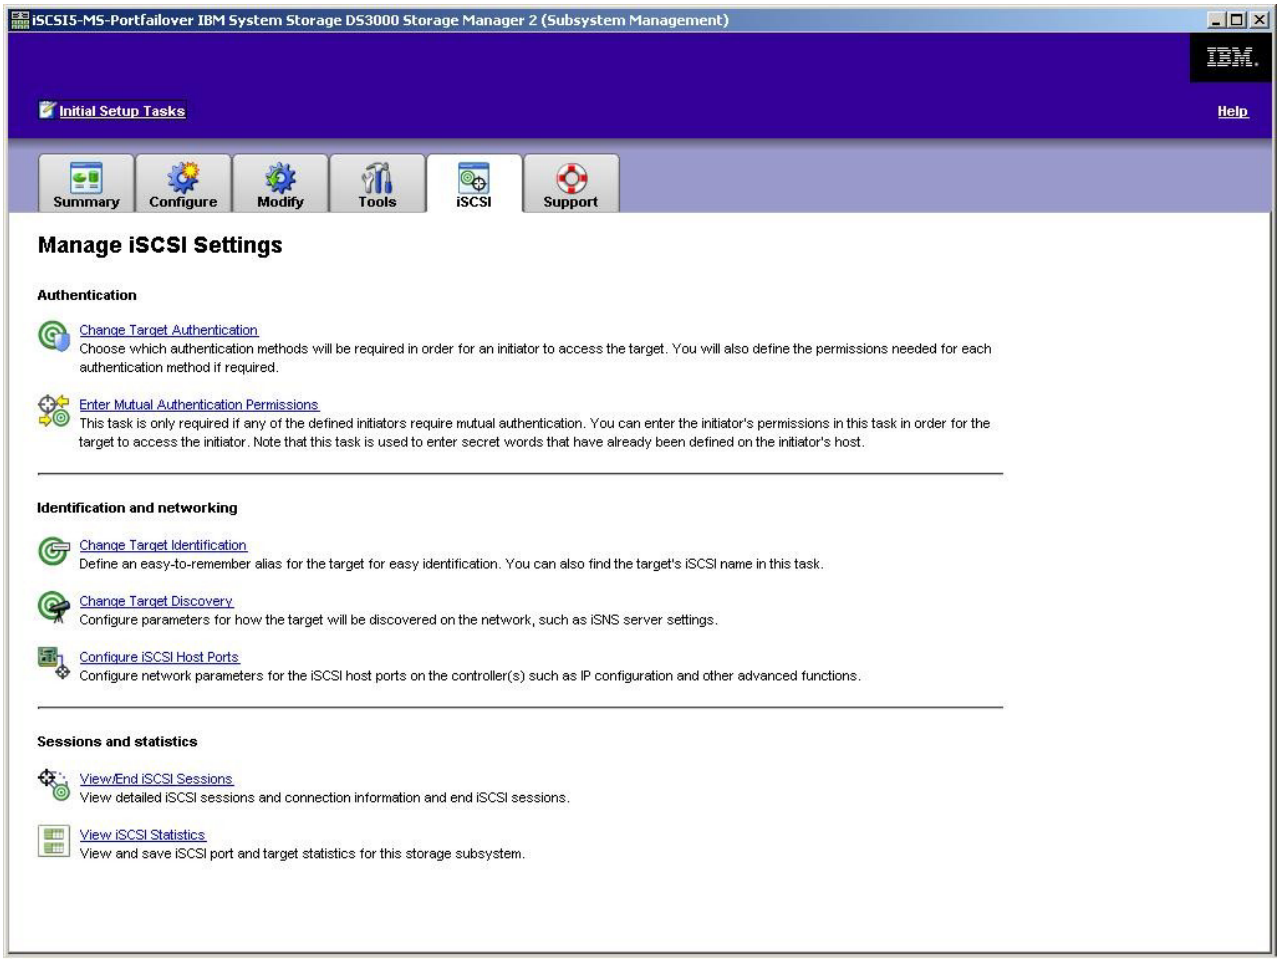

Figure 16. Manage iSCSI settings page

The following options are available from the **iSCSI** page:

- Change Target Authentication
- Enter Mutual Authentication Permissions
- Change Target Identification
- Change Target Discovery
- Configure iSCSI Host Ports
- View/End iSCSI Sessions
- View iSCSI Statistics

Changing target authentication

Select **Change Target Authentication** to specify the target challenge handshake authentication protocol (CHAP) secret that the initiator must use during the security negotiation phase of the iSCSI login. By default, **None** is selected. To change the selection, click **CHAP**, and then type the CHAP secret. You can select the option to generate a random secret if you want to. This enables 1-way CHAP.

Entering mutual authentication permissions

Before you select **Enter Mutual Authentication Permissions**, you must define a host port for the initiator and enable **Target Authentication**. After the host port is listed, select the host from the list and click **Chap Secret** to specify the secret that is passed to the initiator from the target to authenticate it. This enables Mutual CHAP (2-way).

Changing target identification

Select **Change Target Identification** to specify a target alias that is to be used during device discovery. You must provide a unique name for the target that consists of fewer than 30 characters.

Note: You will connect to the target by using the fully qualified IQN that is listed above the alias.

Changing target discovery

Select **Change Target Discovery** to perform device discovery by using the iSCSI simple naming service (iSNS). After you select this option, select the **Use iSNS Server** check box. You can also select whether the iSNS server is discovered using a DHCP server on your network, and you can manually specify an Internet Protocol version 4 (IPv4) or IPv6 address. When you click the **Advanced** tab, you can assign a different TCP/IP port for your iSNS server for additional security.

Note: To provide the required port login information for correct device discovery, all iSCSI ports must be able to communicate with the same iSNS server.

Configuring iSCSI host ports

Select **Configure iSCSI Host Ports** to configure all of the TCP/IP settings. You can choose to enable or disable IPv4 and IPv6 on all of the ports. You can also statically assign IP addresses or let them be discovered using DHCP. Under **Advanced IPv4 Settings**, you can assign VLAN Tags (802.1Q) or set the Ethernet Priority (802.1P). Under **Advanced Host Port Setting**, you can specify a unique iSCSI TCP/IP port for that target port. You can also enable Jumbo Frames from this option. The supported frame sizes are 1500 and 9000.

Viewing or ending an iSCSI session

Select **View/End iSCSI Sessions** to view all of the connected iSCSI sessions to the target. From this page, you can also close an existing session by forcing a target ASYNC logout of the initiator session.

Viewing iSCSI statistics

Select **View iSCSI Statistics** to view a list of all iSCSI session data, for example, the number of header digest errors, number of data digest errors, and successful protocol data unit counts. You can also set a baseline count after a corrective action to determine whether the problem is solved.

Determining firmware levels

There are two methods for determining the DS3000 storage subsystem, expansion unit, drive, and ESM firmware versions. Each method uses the Storage Manager Client that manages the DS3000 storage subsystem with the attached expansion unit.

Method 1

Open the Subsystem Management window and click the **Summary** tab. Select **Storage Subsystem Profile** in the **Hardware Components** area. When the Storage Subsystem Profile window opens, select one of the following tabs to view the firmware information.

Note: The Storage Subsystem Profile window contains all the profile information for the entire subsystem. Therefore, you might have to scroll through a large amount of data to find the information that you want.

Summary

- Firmware version (controller firmware)
- NVSRAM version

The following example shows the profile information on the Summary page.

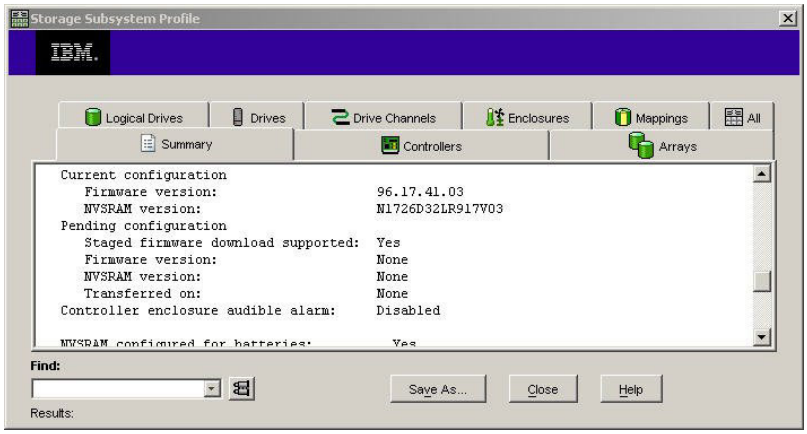

Figure 17. Example of profile information on the Summary page

Controller

- Firmware version
- Appware version (Appware is a reference to controller firmware)
- Bootware version (Bootware is a reference to controller firmware)
- NVSRAM version

The following example shows the profile information on the Controller page.

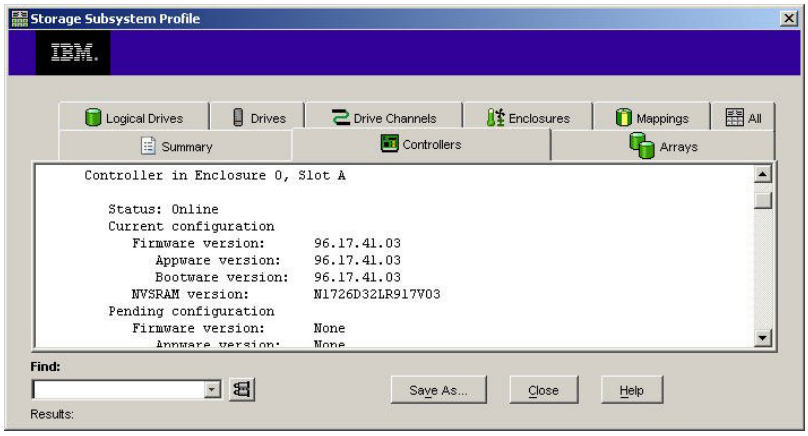

Figure 18. Example of profile information on the Controller page

Enclosures

- Firmware version

The following example shows the profile information on the Enclosures page.

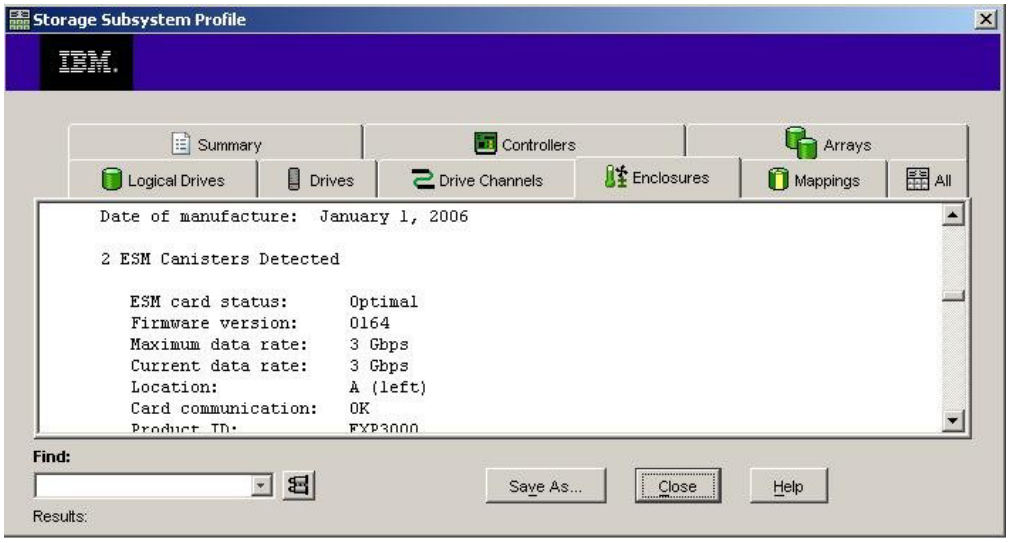

Figure 19. Example of profile information on the Enclosures page

Drives

- Firmware version (drive firmware)

The following example shows the profile information on the Drives page.

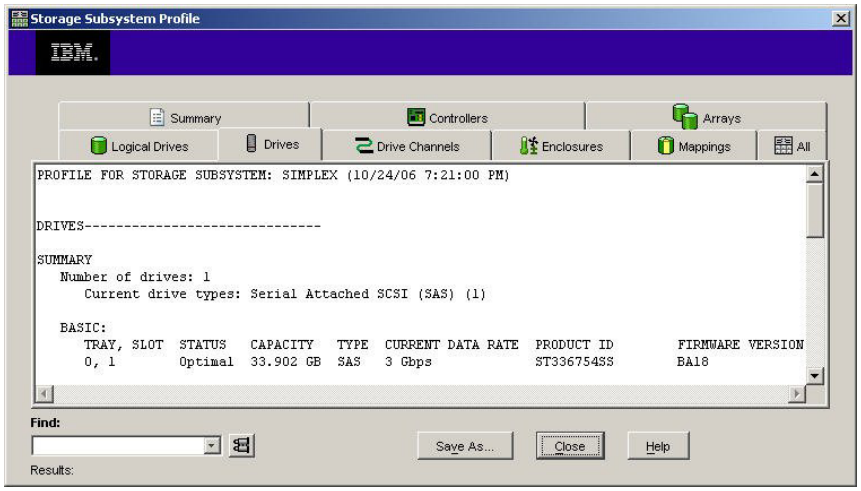

Figure 20. Example of profile information on the Drives page

Method 2

Open the Subsystem Management window and select **Download Firmware** on the Support page. Select one of the following options to view the firmware information.

Note: Use method 2 only to make sure that the firmware versions are correct. After you check the firmware versions, cancel or close the windows to avoid downloading firmware unintentionally.

Download controller firmware

- Firmware version (controller firmware)
- NVSRAM version

Download NVSRAM

- NVSRAM version

Download drive firmware

- Firmware version (drive firmware)

Download Environmental Services Monitor (ESM) firmware

- Enclosure firmware version

Downloading controller, NVSRAM, ESM, and hard disk drive firmware

This section provides instructions for downloading DS3000 storage subsystem controller firmware and NVSRAM, EXP3000 ESM firmware, and drive firmware. The DS3000 storage subsystem firmware download sequence is as follows:

- Controller firmware
- NVSRAM
- ESM firmware
- Hard disk drive firmware

Important: Stop all I/O activity while you download firmware and NVSRAM to a DS3000 storage subsystem, because you will lose the connections between the host server and the DS3000 storage subsystem.

Note: You can download the latest versions of DS3000 storage subsystem controller firmware, NVSRAM, EXP3000 ESM, and hard disk drive firmware at http://www.ibm.com/servers/storage/support/disk/.

Downloading controller or NVSRAM firmware

To download the DS3000 storage subsystem controller firmware and NVSRAM, complete the following steps:

1. From the Enterprise Management window, open the Subsystem Management window by double-clicking a storage subsystem.
2. Click **Support → Download firmware → Download Controller Firmware**. The Download Controller Firmware window opens, as shown in Figure 21.

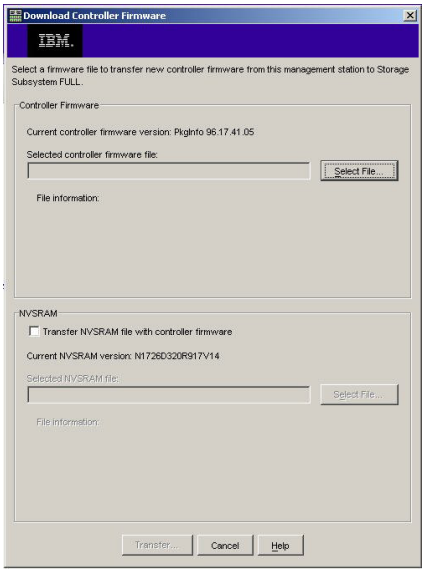

Figure 21. Download Controller Firmware window

3. To select the controller firmware file, click **Select File** next to the **Selected controller firmware file** field, and then browse to the file that you want to download.
4. To select the NVSRAM file, click **Select File** next to the **Selected NVSRAM file** field, and then browse to the file that you want to download.

   To transfer the NVSRAM file with the controller firmware, select the **Transfer NVSRAM file with controller firmware** check box.
5. Click **Transfer**.

Downloading ESM firmware

To download the ESM firmware, complete the following steps.

Note: Downloading ESM firmware must be performed with all I/O quiesed.

1. In the Subsystem Management window, click **Support → Download firmware → Download Environmental (ESM) Card Firmware**. The Download Environmental (ESM) Card Firmware window opens.

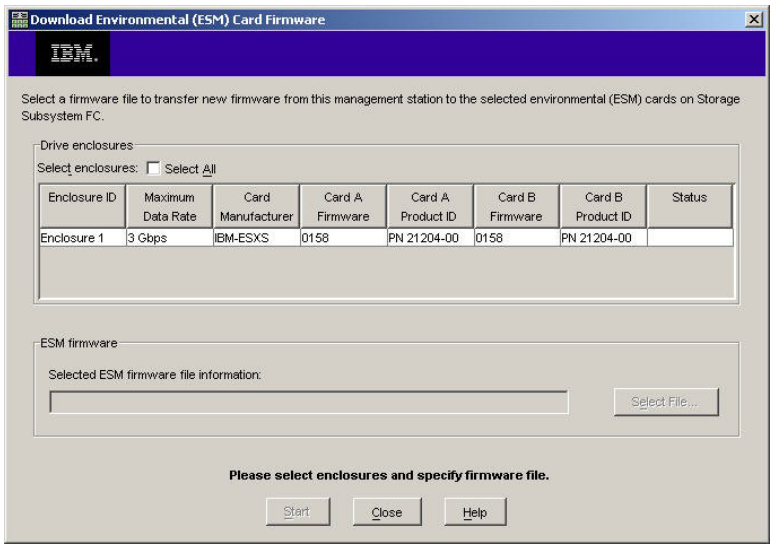

Figure 22. Download Environmental (ESM) Card Firmware window

2. Select the **Select All** check box to direct the download to all enclosures. You can also select one enclosure or combinations of enclosures by pressing the Ctrl key while you select the individual enclosures.
3. To select the ESM firmware file, click **Select File** next to the **Selected ESM firmware file information** field, and then browse to the file that you want to download.
4. Click **Start** to begin the ESM firmware download. A Confirm Download window opens.
5. Type yes and click **OK** to start the download process.
6. When the ESM firmware download for all selected enclosures is completed, click **Close**.

Automatic ESM firmware synchronization

When you install a new ESM in an existing storage expansion enclosure that is connected to a DS3000 storage subsystem that supports automatic ESM firmware synchronization, the firmware in the new ESM is automatically synchronized with the firmware in the existing ESM. This resolves any ESM firmware mismatch conditions automatically.

Note: Automatic ESM firmware synchronization can be performed only after an ESM firmware file is successfully downloaded to an ESM in the storage expansion enclosure.

To enable automatic ESM firmware synchronization, make sure that your system meets the following requirements:

- The Storage Manager Event Monitor is installed and running.
- The DS3000 storage subsystem is defined in the Storage Manager Client (SMclient) Enterprise Management window.

Downloading drive firmware

This section provides instructions for downloading DS3000 drive firmware. For more information, see the online help.

Important: Before you start the drive firmware download process, complete the following tasks:

- Stop all I/O activity before you download drive firmware to a DS3000 storage subsystem.
- Unmount the file systems on all logical drives that access the drives that you select for firmware upgrade.
- Complete a full backup of all data on the drives that you select for firmware upgrade.

To download the drive firmware, complete the following steps:

1. On the Enterprise Management window, open a subsystem management window by double-clicking a storage subsystem.
2. Click **Support → Download firmware → Download Drive Firmware**. The Download Drive Firmware windows opens.

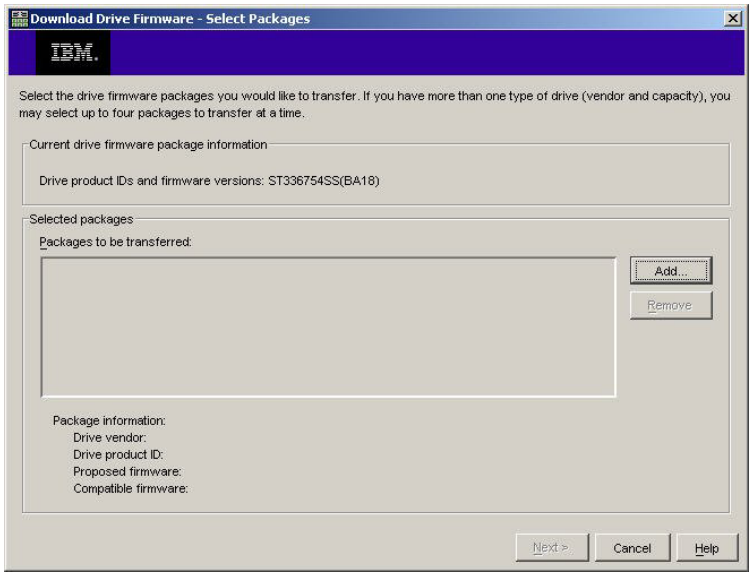

Figure 23. Download Drive Firmware window

3. Click **Add** to locate the server directory that contains the firmware that you want to download.
4. Select the firmware file that you want to download and click **OK**. The file is then listed in the **Selected Packages** area.
5. Select the firmware file for any additional drive types that you want to download, and click **OK**. The additional files are then listed in the **Selected Packages** area. A maximum total of four drive types are possible.

6. Click **Add** to repeat step 5 on page 97 until you have selected each firmware file that you want to download.
7. When you have finished specifying the firmware packages for download, click **Next**.
8. In the Select Drive window (shown in Figure 24), click the **Compatible Drives** tab. The Compatible Drives page contains a list of the drives that are compatible with the firmware package types that you selected. Press and hold the Ctrl key while you use the mouse to select multiple drives individually, or press and hold the Shift key while you use the mouse to select multiple drives that are listed in series. The compatible firmware that you selected in steps 4 on page 97 and 5 on page 97 are downloaded to the drives that you select.

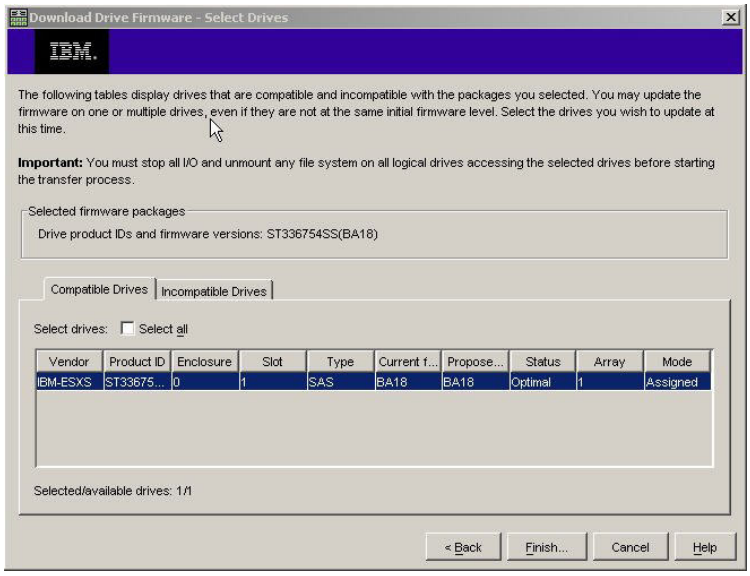

Figure 24. Select Drive window

Note: The firmware packages that you select to download should be listed on the Compatible Drives page. If the product ID for your drive matches the firmware type but it is not listed as compatible on the page, contact your technical-support representative for additional instructions.

9. Click **Finish** to initiate the download of the drive firmware to each compatible drive that you selected in step 8.
10. When the Download Drive Firmware warning window opens and the message `Do you want to continue?` is displayed, type `yes` and click **OK** to start the drive firmware download. The Download Progress window opens, as shown in Figure 25 on page 99. Do not intervene until the download process is completed.

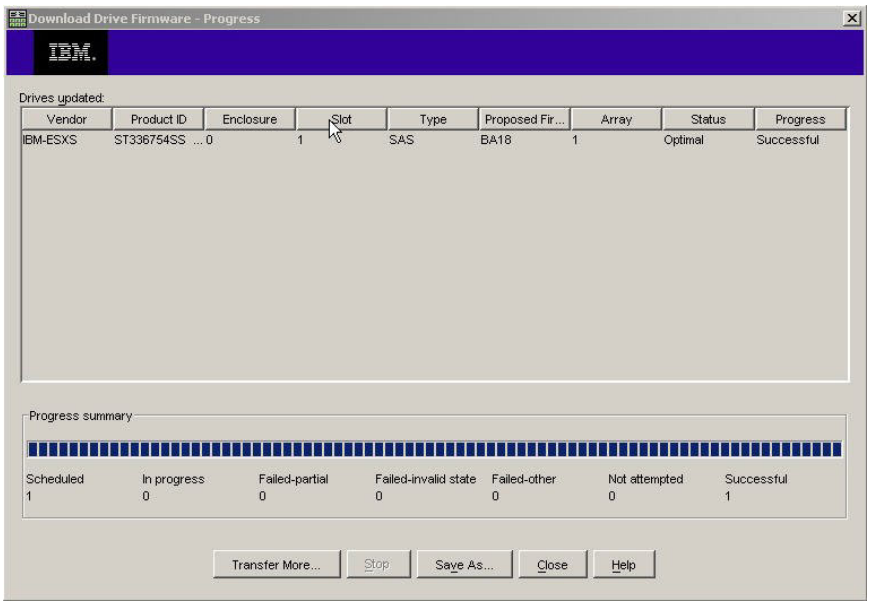

Figure 25. Download Progress window

Every drive that is scheduled for firmware download is designated as `in progress` until it is designated as `successful` or `failed`.

11. If the download for a drive fails, complete the following steps:
    a. Click the **Save as** button to save the error log.
    b. On the Subsystem Management window, click **Support → View Event Log** and complete the following steps to save the storage subsystem event log before you contact your technical-support representative:
        1) Click **Select all**.
        2) Click **Save as**.
        3) Provide a file name to which to save the log.

    When the **Close** button is active, the drive firmware download process is completed.
12. Click **Close**.

Array import and export functions

The array import and export functions enable the transfer of configurations between DS3000 storage subsystems.

Important: The array import and export functions are available for only DS3000 storage subsystems that are running controller firmware version 07.35 or later. This function does not work with storage subsystems or arrays that are using controller firmware version 06.xx.

Exporting an array

To export an array, complete the following steps:

1. In the Subsystem Management window, click the **Advanced Support** tab.
2. Click **Export Array**. The Export Array wizard opens. The wizard helps you prepare an array for export from one storage subsystem to another storage subsystem.

Important: Before you start the export function, you must stop all I/O to the array that you are exporting.

3. Select the array that you want to export and click **Next**.

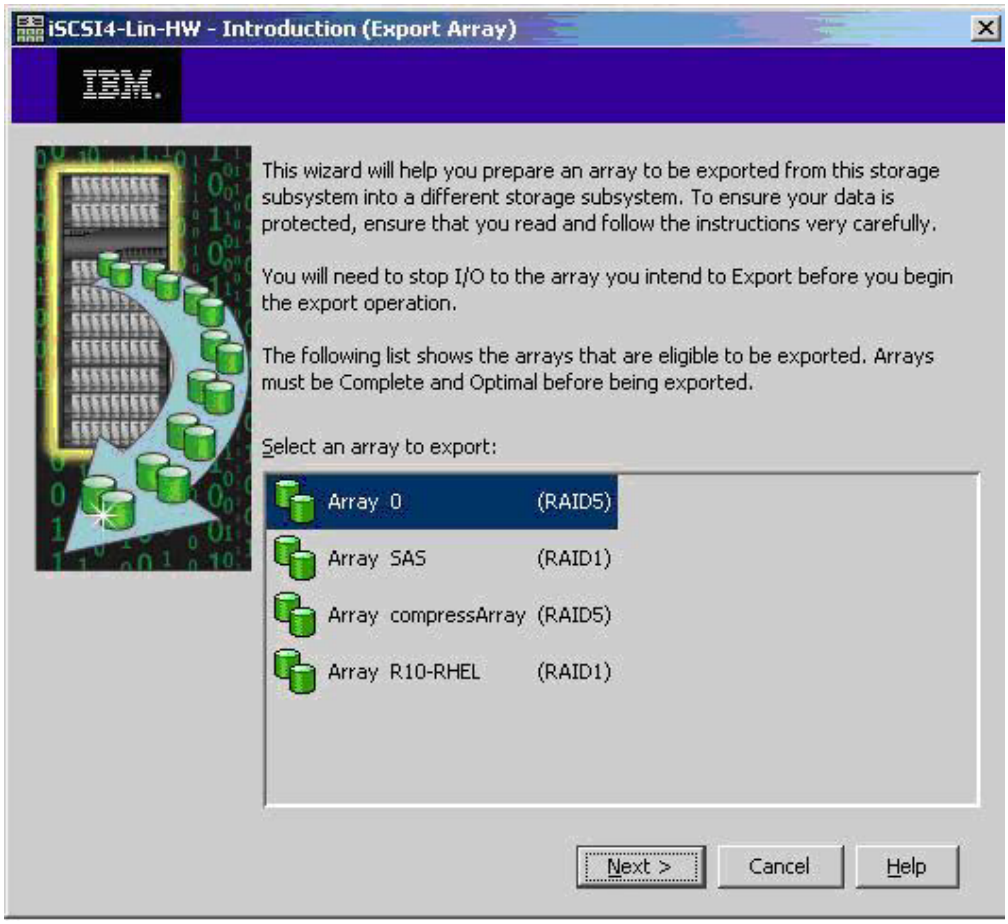

Figure 26. Selecting an array in the Export Array wizard

4. The Preparation Checklist window opens and provides a list of procedures that you must complete before you can export an array.

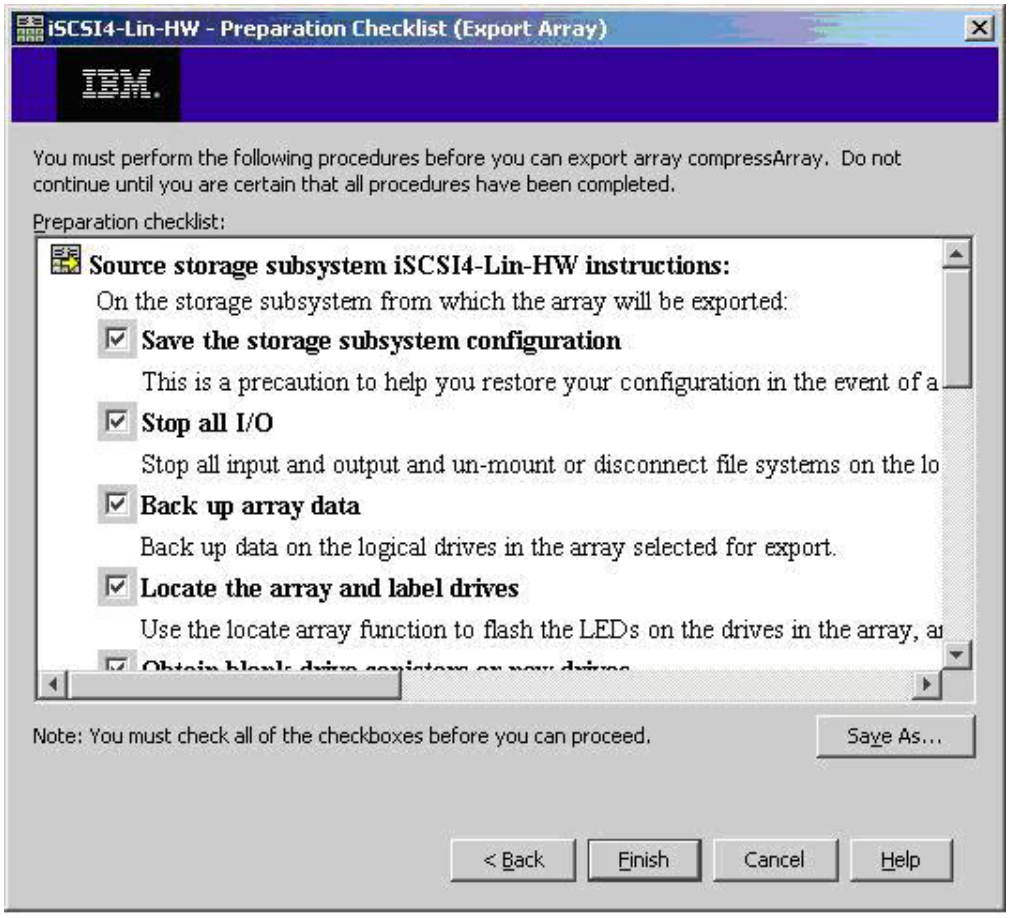

Figure 27. The Export Array Preparation Checklist window

5. After you complete the preparation checklist, click **Finish**.
6. In the Confirm Export window, type yes to begin the import operation process, and then click **OK**.
7. When the export operation is completed, you can remove the hard disk drives from the storage subsystem.

Importing an array

To import an array, complete the following steps:

1. Insert the hard disk drives that are associated with the exported array in the storage subsystem.
2. In the Subsystem Management window, click the **Advanced Support** tab.

3. Click **Import Array**. The Import Report window opens and shows the details of the array that you are importing.

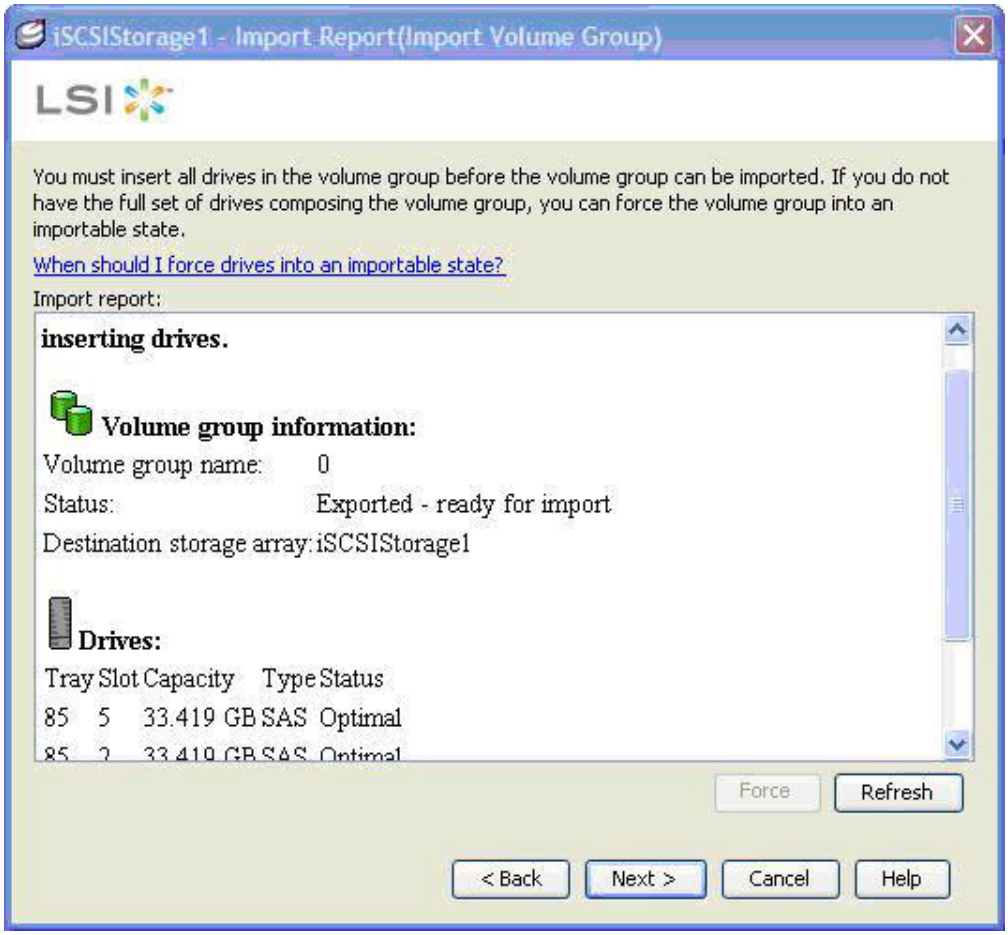

Figure 28. Import Report window

4. If the information in the report is correct, type yes in the Confirm Import window, and then click **OK** to begin the import operation process.

Performing other storage-subsystem management tasks

The following list describes other storage-subsystem management tasks that you can perform. You can perform the following tasks from the Subsystem Management window.

- Create hot spares (**Configure → Configure Hot Spares**)
- Automatically create logical drives and hot spares (**Configure → Automatic Configuration**)
- Edit host topology configurations for the subsystem (**Modify → Edit Topology**)
- Enter or change a storage subsystem password (**Tools → Set or Change Password**)
- View or change media scan settings (**Tools → Change Media Scan Settings**)
- Locate a storage subsystem or its components (**Tools → Locate**)
- View or enable premium features (**Tools → View/Enable Premium Features**)
- Change enclosure identification numbers for a storage subsystem (**Tools → Change Enclosure ID Numbers**)

- Synchronize controller clocks (**Tools → Synchronize Controller Clocks**)
- Change the network configuration (**Tools → Change Network Configuration**)
- Reset the battery age (**Tools → Reset Battery Age**)
- Inherit system settings from the operating system (**Tools → Inherit System Settings**)
- View the major event log (**Support → View Event Log**)
- Save all storage subsystem data (**Support → Gather Support Information**)
- Place controller online or offline (**Support → Manage Controllers**)

For more information about these and other storage-subsystem management tasks, see the applicable topics in the Subsystem Management online help.

For advanced users: You can perform other storage-subsystem management tasks from the Enterprise Management window by clicking **Tools → Script Editor**. For more information, see the Enterprise Management window online help.

Note: If there is a problem with the storage subsystem, a Needs Attention status icon is displayed in the Enterprise Management window next to the storage subsystem. In the Subsystem Management window, a Storage Subsystem Needs Attention link is displayed in the status area of the Summary page. You can click the link to open the Recovery Guru. You can also open the Recovery Guru by clicking **Support → Recover from Failure** in the Subsystem Management window.

Best practices guidelines for the DS3300

For the best performance of the Storage Manager software and the DS3300 , follow the guidelines in the following sections.

iSNS best practices

There are many considerations for using an iSNS server correctly. Make sure that you correctly assign your iSNS server address that is provided during the DHCP lease discovery of your initiator or target. This enables ease of discovery when you use software initiator-based solutions. If you are unable to do this and must manually assign the iSNS server to your software or hardware initiators, you should make sure that all ports for both the DS3300 target and the iSCSI initiator are in the same network segment (or make sure that the routing between the separate network segments is correct). If you do not do this, you will be unable to discover all ports during the iSCSI discovery process, and you might not be able to correctly perform a controller or path failover.

Using DHCP

Using DHCP for the target portals is not recommended. If you use DHCP, you should assign DHCP reservations so that leases are maintained consistently across restarts of the DS3300 storage subsystem. If static IP reservations are not provided, the initiator ports can lose communication to the DS3300 controller and might not be able to reconnect to the device.

Using supported hardware initiators

As of the date of this document, the only supported hardware initiator is the QLogic iSCSI Expansion Card for IBM @server BladeCenter.

All of the hardware initiators that are supported use the same base firmware code and the SANsurfer management application. Before you install and configure these adapters, make sure that you have installed the latest management application and the latest firmware code. After you confirm this, configure each adapter one at a time.

To make sure that failovers are performed correctly, connect each adapter by using one of the following two basic configurations:

- If you have a simple configuration in which all adapters and target ports are in the same network segment, each adapter should be able to log in to any target port.
- If you have a complex configuration, each adapter is allowed a single path to each controller device.

To log in correctly to all available target ports from the hardware initiator, complete the following steps.

Note: Failure to perform the steps in the following procedure might result in path failover inconsistencies and incorrect operation of the DS3300.

1. Start the SANsurfer management utility.
2. Connect to the system that is running the qlremote agent.
3. Select the adapter that you want to configure.
4. Select Port 0 or Port 1 for the adapter.
5. Click **Target Settings**.
6. Click the plus sign (+) in the far right of the window.
7. Type either the IPv4 or IPv6 address of the target port to which you want to connect.
8. Click **OK**.
9. Select **Config Parameters**.
10. Scroll until you see **ISID**.
11. For connection 0, the last character that is listed should be 0. For connection 1, it should be 1, for connection 2, it should be 2, and so on.
12. Repeat steps 6 through 11 for each connection to the target that you want to create.
13. After all of the sessions are connected, select **Save Target Settings**.

If you are using the QLogic iSCSI Single-Port or Dual-Port PCIe HBA for IBM System x™ to support IPv6, you should allow the host bus adapter firmware to assign the local link address.

Using IPv6

The DS3300 supports the Internet Protocol version 6 (IPv6) TCP\IP. Note that only the final four octets can be configured if you are manually assigning the local link address. The leading four octets are fe80:0:0:0. The full IPv6 address is required when you are attempting to connect to the target from an initiator. If you do not provide the full IPv6 address, the initiator might fail to be connected.

Network settings

Using the DS3300 in a complex network topology introduces many challenges. If possible, try to isolate the iSCSI traffic to a dedicated network. If this is not possible, follow these suggestions:

- If you are using a hardware-based initiator, the Keep Alive timeout should be 120 seconds. To set the Keep Alive timeout, complete the following steps:
  1. Start the SANsurfer Management Utility and connect to the server.
  2. Select the adapter and the adapter port that is to be configured.
  3. Select the port options and firmware.

     The default connection timeout is 60 seconds. This setting is correct for simple network topologies; however, if a network convergence occurs and you are not using Fast Spanning Tree and separate spanning tree domains, in a more complex configuration, you might encounter I/O timeouts.
- If you are using a Linux software initiator to connect to the DS3300, modify the ConnFailTimeout to account for the spanning tree issue that is described in step 3. The ConnFailTimeout value should be set to 120 seconds.

Operating system consideration

For the best performance, do not enable Data and Header Digest when using the embedded Software Initiators with the Red Hat Enterprise Linux 4 and SUSE Linux Enterprise Server 9 operating systems. If these functions are enabled, you will see degraded performance and in the case of multiple hosts accessing the same DS3300 storage subsystem, you might see that a path is marked incorrectly as having failed.

SATA disk drive best practices guidelines

If you use a stripe size of 8 KB, you cannot perform a format of a Linux ext2 file system on a logical drive configured with SATA drives. For the best performance, use a stripe size of 32 KB or larger when using SATA disk drives.

Fibre Channel best practices guidelines

If you install your Linux host with the QLogic host bus adapter installed in the host, the default device driver provided with the distribution is installed. This device driver has embedded failover support enabled that prevents the correct discovery of all available controller paths. You must install the most current HBA device drivers that are provided for your solution and make sure that the ramdisk image is remade with those device drivers before you install the MPP failover driver. Failure to do so results in incorrect controller and path failover.

Storage Manager 2 Enterprise Manager view and Subsystem Management view best practices guidelines

Consider the following information before you use the Storage Manager 2 Enterprise Management view and Application Management view:

- The current command to gather performance statistics (save storagesubystem performancestats=filename.xls;) does not gather any valid data. All of the values are zero (0). As of the date of this document, there is no workaround available.
- To create a logical drive using the smcli or Enterprise Manager script editor you must make nsure that the syntax of the command is correct. When you specify a logical drive size, the correct syntax is:

  `create logicalDrive array[`*`array_number`*`] capacity=`*`size_of_logical_drive TB/GB/MB`*`;`

  Note: There must be a space before *TB/GB/MB*; otherwise, the command does not work.

- The Storage Manager 2 Application Management view is limited to creating a maximum of 64 FlashCopies and 128 VolumeCopies. However, using the Storage Manager command-line interface enables a maximum of 128 FlashCopies and 255 VolumeCopies to be created.
- If you are updating a DS3000 storage subsystem, the first controller firmware download might fail. If this occurs, start the firmware download a second time and verify that the firmware is successfully downloaded. If the firmware download fails again, contact your IBM technical-support representative.
- To expand a logical drive using the smcli or Enterprise Manager script editor you must make nsure that the syntax of the command is correct. When you specify a logical drive size, the correct syntax is:

  ```
  set logicalDrive[logical_drive_name]
  addCapacity=size_to_increase_logical_drive_by TB/GB/MB;
  ```

  Note: There must be a space before *TB/GB/MB*; otherwise, the command does not work.
- If you are using a firewall between your management client and the host agent that is managing the controller or the controller TCP/IP address, make sure that you have created an exclusion for TCP/IP port 2463. This port is used by the Storage Manager software to communicate to the storage subsystem.
- If you perform an upgrade to the Storage Manager software from any previous version to Storage Manager version 02.70.xx.xx, you must perform a new discovery of all your storage subsystems again. The previous management domain is not maintained.

Chapter 6. Completing storage-management tasks

This chapter provides information about how to use the SMdevices utilities, start and stop the host-agent software, and uninstall the Storage Manager software.

Using the SMdevices utility

The SMutil software includes a utility, SMdevices, that you can use to view the storage subsystem logical drive that is associated with an operating-system device name. This utility is helpful when you want to map host system LUNs with DS3000 LUNs.

Note: The SMdevices command does not function with the DS3300 storage subsystem.

When you are finished creating the logical drives on a storage subsystem, to use SMdevices, complete the following steps on the host that is attached to that storage subsystem:

1. At a command prompt, change to the util directory, where the IBM DS3000 Storage Manager software is installed on your host system.
2. Enter the following command:

   ```
   # SMdevices
   ```

   The software displays the following device identification information:

   ```
     /dev/sdb (/dev/sg1) [Storage Subsystem CGG_SATA_Test, Logical Drive data-1,
   LUN 0, Logical Drive ID <600a0b80001d2b9500004d80469bbb18>,
   Preferred Path (Controller-A): In Use]
   ```

   In this example:

   - */dev/sdb (/dev/sg1)* is the disk number that is displayed in Disk Administrator
   - *CGG_SATA_Test* is the name of the DS3000 storage subsystem
   - *data-1* is the name of the logical drive
   - *LUN 0* is the LUN that is associated with the logical drive

Stopping and restarting the host-agent software

You must stop and restart the host-agent software if you add storage subsystems to the management domain of the host-agent software. When you restart the service, the host-agent software discovers the new storage subsystems and adds them to the management domain.

Note: If none of the access logical drives are detected after a restart, the host-agent software automatically stops running. Make sure that the Fibre Channel or iSCSI connection from the host to the SAN to which the DS3300 or DS3400 storage subsystem is attached is working correctly. Then, restart the host or cluster node so that new host-agent-managed storage subsystems can be discovered.

To stop and restart the host-agent software, use the following commands:

`# SMagent start` starts the SMagent

`# SMagent stop` stops the SMagent

Uninstalling storage management software components

Use the following procedure to uninstall one or more of the components of the Storage Manager 2 software. Existing storage array mappings and storage partition configurations are retained during software removal and will be recognized by the new client software.

Important: Do not uninstall the MPIO PCM or RDAC MPP device driver unless you are instructed to do so by your technical-support representative.

To uninstall the Storage Manager 2 software, complete the following steps:

1. Open a terminal window and change to the Uninstall IBM_DS3000 directory.
2. Enter the following command:

   ```
   # sh Uninstall_IBM_DS3000
   ```

3. Choose one of the following uninstall options:
   a. Completely remove all features and components
   b. Choose specific features that were installed
4. Select option **a** and press Enter to complete the uninstall process.
5. Select option **b** and then select from the following features to uninstall:
   - Storage Manager 2 Client
   - Storage Manager 2 Utilities
   - Storage Manager 2 Agent

   Enter the number of the feature that you want to uninstall and press Enter to complete the uninstall process.

Chapter 7. Enabling and using premium features

Premium features provide enhanced functionality that enables you to perform tasks that might be limited or unavailable with your base storage subsystem model.

You can purchase the following premium features for a DS3000 storage subsystem:

- DS3000 FlashCopy Expansion License
- DS3000 Volume Copy License
- DS3000 FlashCopy Volume Copy License
- DS3000 Partition Expansion License

For information about how to purchase a premium feature, contact your IBM marketing representative or authorized reseller.

Obtaining the activation key file and enabling a premium feature

Follow the *Activation Instructions* document that comes with your premium feature to obtain the activation key file. The activation key file is required by the Storage Manager software to enable the premium feature.

Enabling the DS3000 Partition Expansion License

You can upgrade the DS3000 Partition Expansion License from 4 to 16 partitions for the storage subsystem. After you have obtained the DS3000 Partition Expansion License key, to enable this feature in the Storage Manager software, complete the following steps:

1. Open the Subsystem Management window and click **Tools → View/Enable Premium Features**.
2. Under **Enabled Premium Features**, click **Upgrade a feature**, as shown in Figure 29 on page 110.

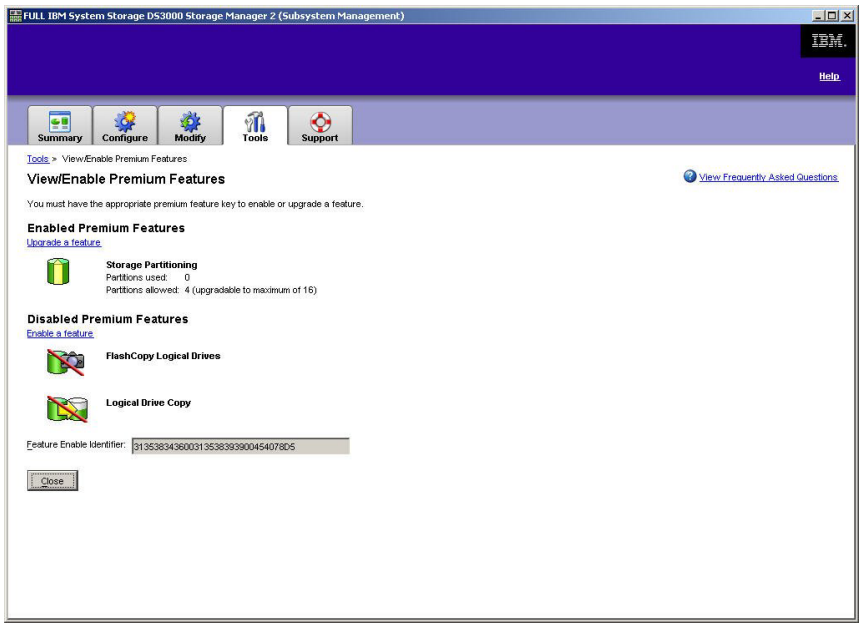

Figure 29. View/Enable Premium Features window

The Select Feature Key File window opens.

3. Select the feature key file that you obtained from the IBM Web site and click **OK**.
4. In the Enable Premium Feature window, click **Yes**.

Enabling the DS3000 FlashCopy Expansion License

After you have obtained the DS3000 FlashCopy Expansion License key, to enable this feature in the Storage Manager software, complete the following steps:

1. Open the Subsystem Management window and click **Tools → View/Enable Premium Features**.
2. Under **Disabled Premium Features**, click **Enable a feature**, as shown in Figure 29. The Select Feature Key File window opens.
3. Select the feature key file that you obtained from the IBM Web site and click **OK**.
4. In the Enable Premium Feature window, click **Yes**.

Enabling the DS3000 VolumeCopy License

Note: The terms *VolumeCopy* and *Logical Drive Copy* are used interchangeably throughout this document, the Storage Manager Client program interface, and the online help.

To enable the DS3000 Volume Copy License, complete the following steps:

1. Open the Subsystem Management window and click **Tools → View/Enable Premium Features**.
2. Under **Disabled Premium Features**, click **Enable a feature**, as shown in Figure 29. The Select Feature Key File window opens.
3. Select the feature key file that you obtained from the IBM Web site and click **OK**.

4. In the Enable Premium Feature window, click **Yes**.

Using the FlashCopy premium feature

A FlashCopy logical drive is a logical point-in-time image of a logical drive, called a base logical drive. A FlashCopy logical drive has the following features:

- It is created quickly and requires less disk space than an actual logical drive.
- It can be assigned a host address, so that you can perform backups by using the FlashCopy logical drive while the base logical drive is online and accessible.
- You can use the FlashCopy logical drive to perform application testing or both scenario development and analysis. This does not affect the actual production environment.
- The maximum number of FlashCopy logical drives that is allowed is one-half of the total logical drives that are supported by your controller model.

For more information about the FlashCopy feature and how to manage FlashCopy logical drives, see the Subsystem Management online help.

Important: The FlashCopy drive cannot be added or mapped to the same server that has the base logical drive of the FlashCopy logical drive in a Windows Server 2003 or NetWare environment. You must map the FlashCopy logical drive to another server.

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

1. To make sure that you have the accurate point-in-time image of the base logical drive, stop applications and flush cache I/O to the base logical drive.
2. Open the Subsystem Management window, click **Configure → Create FlashCopy Logical Drive**, and follow the instructions in the wizard.
3. For instructions for adding the FlashCopy logical drive to the host, see the Subsystem Management online help

Note: To manage created FlashCopy Logical Drives, click the **Modify** tab and select **Modify FlashCopy Logical Drives**. Select **Disable FlashCopy Logical Drives**, **Re-create FlashCopy Logical Drives**, or **Expand the FlashCopy Repository** and then follow the instructions in the wizard.

Using VolumeCopy

The VolumeCopy feature is a firmware-based mechanism for replicating logical drive data within a storage array. This feature is designed as a systems-management tool for tasks such as relocating data to other drives for hardware upgrades or performance management, data backup, or restoring snapshot volume data. You submit VolumeCopy requests by specifying two compatible drives. One drive is designated as the source and the other as the target. The VolumeCopy request is persistent so that any relevant result of the copy process can be communicated to you.

For more information about the VolumeCopy feature and how to manage VolumeCopy logical drives, see the Subsystem Management online help.

To create a VolumeCopy, open the Subsystem Management window, click **Configure → Create Volume Copy**, and follow the instructions in the wizard

Note: To manage created VolumeCopy logical drives, click **Modify → Manage Logical Drive Copies**. A list of VolumeCopy requests is displayed. For each

VolumeCopy, select **Re-Copy** or **Stop**, and to modify attributes, select **Permissions** and **Priority**. Figure 30 shows this tool.

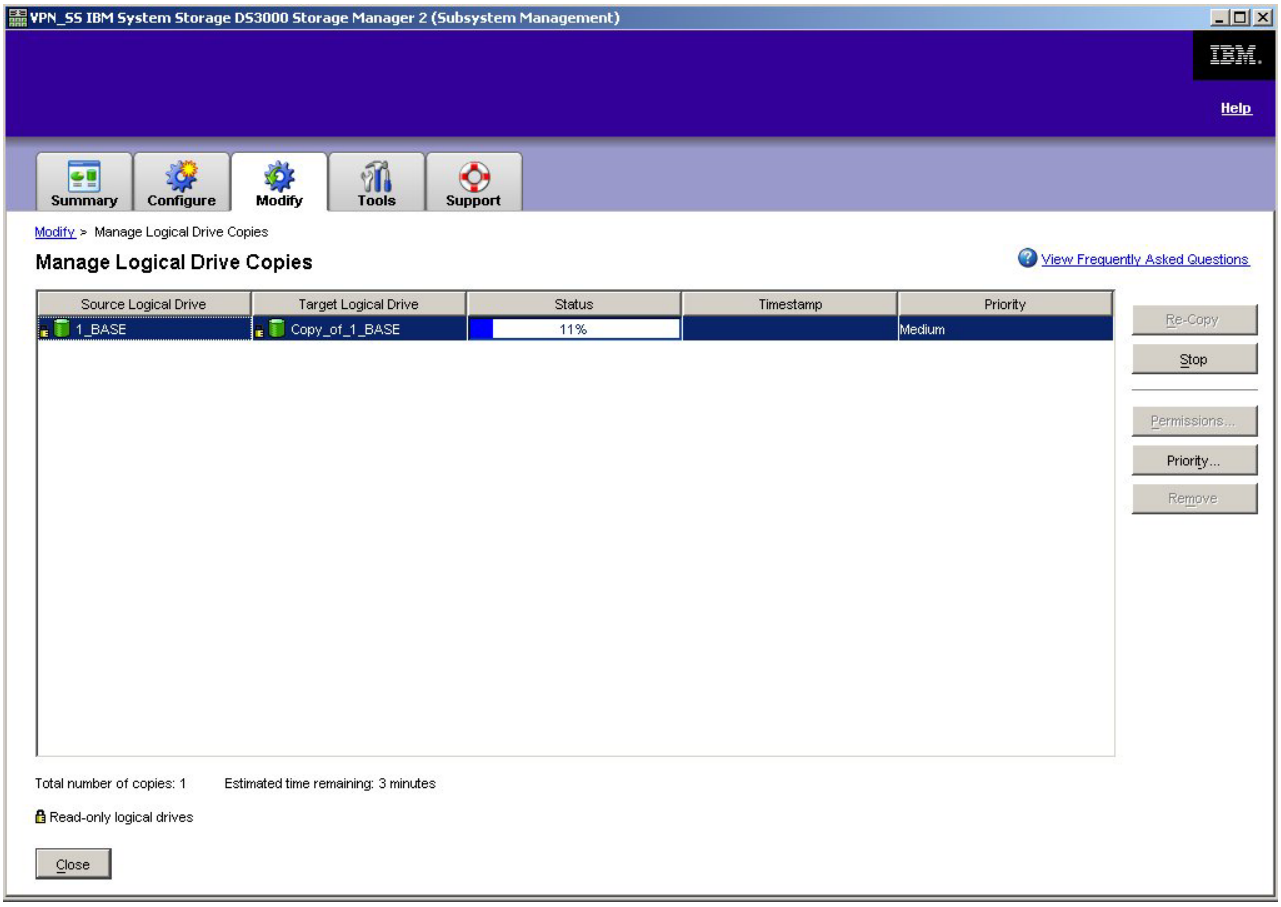

Figure 30. Manage Logical Drive Copies window

Appendix A. Storage subsystem and controller information record

Table 9 provides a data sheet for recording storage subsystem names, management types, Ethernet hardware addresses, and IP addresses. Make a copy of this table and complete the information for your storage subsystems and controllers. Use the information to set up the BOOTP table for the network server and the host or Domain Name System (DNS) table. The information is also helpful if you add storage subsystems after initial installation. Each column heading includes a page reference for detailed instructions for obtaining the information. For a sample information record, see Table 8 on page 18.

Table 9. Storage subsystem and controller information record

Storage subsystem name (page 18)	**Management method (page 6)**	**Controllers: Ethernet and IP addresses, and host name (pages 18 and 19)**		**Host: IP address and host name (page 19)**

Appendix B. Using the IBM System Storage DS3000 Controller Firmware Upgrade Tool

Important:

Use the IBM System Storage DS3000 Controller Firmware Upgrade Tool only when you are migrating DS3000 controllers from version 06.22 or 06.24 to version 07.35.

Do not use the controller firmware upgrade tool to perform standard upgrades for controller, ESM, or hard disk drive firmware. To perform a standard firmware upgrade for a DS3000 controller, EXP3000 ESM, or a hard disk drive, see “Downloading controller, NVSRAM, ESM, and hard disk drive firmware” on page 94.

Overview

Attention: To prevent the loss of date, before you use the controller firmware upgrade tool to migrate from firmware version 06.22 or 06.24 to version 07.35, you must backup all data and save the existing system configuration. After the upgrade tool has completed a firmware upgrade, the DS3000 controllers cannot be returned to previous firmware version levels.

Before you use the upgrade tool, read the following important information:

- Perform the upgrade offline.
- Perform the overall installation of the Storage Manager software into an existing host environment online.
- For most failover drivers to take effect, restart the host.
- Make sure that all devices have an Optimal status before you download firmware.
- Check the current firmware level to make sure that the firmware level is either 06.22.xx.xx or 06.24.xx.xx. The firmware upgrade tool supports upgrading to 07.xx level firmware from only firmware versions 06.22.xx.xx or 06.24.xx.xx.

Attention: To avoid a potential loss of data access, check the Storage Manager readme file to make sure that the firmware you download is compatible with the Storage Manager software that is installed on your storage subsystem. If non-compatible firmware is downloaded, you might lose access to the drives in the storage subsystem, so upgrade the Storage Manager software first. Do not make changes to your configuration or remove drives or enclosures during the upgrade process.

Checking the storage subsystem health conditions

To determine the health condition of the storage subsystem, complete the following steps:

1. From the Array Management window in the Storage Manager software, right-click the storage subsystem. The Storage Manager software establishes communication with each managed device and determines the current status. There are six possible status conditions:

   Optimal

   Every component in the storage subsystem is in the working condition that you want.

Needs Attention
A problem exists with the storage subsystem that requires intervention to correct it.

Fixing A Needs Attention condition is corrected, and the storage subsystem is currently changing to an Optimal status.

Unresponsive
The management station cannot communicate with one or both controllers in the storage subsystem.

Contacting Device
The Storage Manager software is establishing contact with the storage subsystem.

Needs Upgrade
The storage subsystem is running a level of firmware that is no longer supported by the Storage Manager software.

2. If the status is Needs Attention, write down the condition. Contact your IBM technical-support representative for fault resolution.

   Note: The Recovery Guru in the Storage Manager software also provides a detailed explanation of the conditions and recovery procedures.

Installing the controller firmware upgrade tool

To install the controller firmware upgrade tool, complete the following steps:

1. Open the DS3000 Storage Manager 10 installation file according to your operating system procedures. The IBM System Storage DS3000 Controller Firmware Upgrade Tool is installed as part of the Storage Manager 10 installation.
2. Click **Next**.
3. Accept the license agreement, and click **Next**.
4. Select the folder in which you want to install the tool, and click **Next**.
5. Click **Install**.
6. Click **Done**.

Adding a storage subsystem

To add a storage subsystem using the upgrade tool, complete the following steps:

1. Click **Add**. The Select Addition Method window opens.
2. Select either **Automatic** or **Manual**.
3. Click **OK** to begin adding storage subsystems.
4. To view any issues with the storage subsystem that you added that might impede upgrading the firmware, click **View Log**.

Downloading the firmware

To download the firmware, complete the following steps:

1. Select the storage subsystem for which you want to download the firmware, and then click **Download Firmware**. The Download Firmware window opens.
2. To choose the controller firmware file that you want to download from a directory on your computer or on your network, click **Browse**.

3. To choose the NVSRAM file that you want to download from a directory on your computer or on your network, click **Browse**.
4. Click **OK**. The firmware starts to download. A status bar is displayed in the Controller Firmware Upgrade window.

After the firmware file is downloaded, the firmware begins the activation process, and the status Activating is displayed in the status bar in the controller firmware upgrade tool.

Note: The firmware activation process can take up to 30 minutes to be completed. If after 30 minutes the Activating status message is not changed, check the controller profile data using the Storage Manager 10 software to determine if the firmware download has completed successfully.

Viewing the controller firmware upgrade tool log file

If you encounter any problems updating the controller firmware, complete the following steps to view the controller firmware upgrade tool log file:

1. Click **View Log**. The View Log window opens. This log documents any issues with the storage subsystem that might prevent you from updating the firmware.
2. If any issues are documented in the log, correct those issues before you try to download the firmware.

Appendix C. Solaris failover drivers

A 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.

Solaris host systems require the Solaris Multiplexed I/O (MPxIO) failover driver.

Installing the MPxIO driver

Multiplexed I/O (MPxIO) is a Sun Solaris multipath driver architecture. This failover driver enables storage arrays to be accessed through multiple host controller interfaces from a single instance of the storage array. MPxIO helps protect against storage subsystem outages because of controller failures. If one controller fails, MPxIO automatically switches to an alternate controller. MPxIO is fully integrated within the Solaris 10 operating system.

For more information, go to http://sun.com/docs and see the following Sun documents:

- *Sun Solaris Fibre Channel and Storage Multipathing Administration Guide*
- *Setting Up Solaris iSCSI Multipathed Devices*

Device name change considerations for MPxIO

In the /dev and /devices trees, devices are named differently from their original names when MPxIO is enabled. For example:

Device name with MPxIO disabled

```
/dev/dsk/c1t1d0s0
```

MPxIO-enabled device name

```
/dev/rdsk/c0t600A0B800011121800006B31452CC6A0d0s2
```

You must configure applications that directly consume the device to use the new names whenever the MPxIO configuration is enabled or disabled.

In addition, the /etc/vfstab file and the dump configuration also contain references to device names. When you use the **stmsboot** command to enable or disable MPxIO, as described in the next sections, /etc/vfstab and the dump configuration are automatically updated with the new device names.

Downloading the latest MPxIO driver version

MPxIO is fully integrated within the Solaris 10 operating system, and does not need to be installed separately. MPxIO with Solaris 10 is updated using regular Solaris 10 patches, which are available at http://sun.com/sunsolve.

Note: Install the regular kernel jumbo patch, because there are dependencies between the various patches that make up the driver stack.

Enabling the MPxIO failover driver

You can enable MPxIO by using the **stmsboot** command. The **stmsboot** command also updates the device names in the /etc/vfstab file and the dump configuration files during the next restart.

Note: In Solaris 10, the **stmsboot** command is used to enable or disable MPxIO on all devices.

Before you begin:

1. Install the Solaris operating system and the latest patches.
2. When the host is defined, make sure that you select **Solaris** as the host type.
3. Keep in mind the following considerations for **stmsboot -e [enable] -d [disable] and -u [update]**:
   - When you run the **stmsboot** command, you must select **Reboot the system now** (the default).
   - The **stmsboot** command saves copies of the original /kernel/drv/fp.conf and /etc/vfstab files before modifying them, so you can use the saved files to recover from any unexpected problems.
   - Make sure that the eeprom boot device is set to start from the current boot device.

To enable MPxIO on all I/O devices, complete the following steps:

1. Run the **stmsboot -e** command, and select **y** (the default) to restart the system:

```
# stmsboot -e

WARNING: This operation will require a reboot.
Do you want to continue ? [y/n] (default: y) y
The changes will come into effect after rebooting the system.
Reboot the system now ? [y/n] (default: y) y
```

   Note: During the restartt, /etc/vfstab and the dump configuration are updated to reflect the device name changes.
2. After the reboot, configure your applications to use new device names, as explained in “Device name change considerations for MPxIO” on page 119.
3. If necessary, edit the /kernel/drv/fp.conf configuration file to make sure that the following parameter is set:

   ```
   mpxio-disable="no";
   ```

   Edit the /kernel/drv/scsi_vhci.conf configuration file to make sure that the following parameters are set:

   ```
   load-balance="none";
   auto-failback="enable";
   ```
4. If you made any changes to configuration files in the previous step, save the file, and restart the server by typing the following command:

   ```
   # shutdown  -g0  -y  -i6
   ```
5. If needed, update the HBA firmware.
6. Create the DS3000 logical drives and map them to the host ports in the Sun servers.

Verifying devices and configuring failover / failback path for the mapped LUNs

To verify devices and configure the failover path for the mapped LUNs, complete the following steps:

1. Verify devices using the **cfgadm –al** command to display information about the host ports and their attached devices:

```
# cfgadm -al
Ap_Id                          Type         Receptacle   Occupant     Condition
PCI0                           vgs8514/hp   connected    configured   ok
PCI1                           unknown      empty        unconfigured unknown
PCI2                           unknown      empty        unconfigured unknown
PCI3                           mult/hp      connected    configured   ok
PCI4                           unknown      empty        unconfigured unknown
PCI5                           unknown      empty        unconfigured unknown
PCI6                           unknown      empty        unconfigured unknown
PCI7                           mult/hp      connected    configured   ok
PCI8                           mult/hp      connected    configured   ok
c0                             scsi-bus     connected    configured   unknown
c0::dsk/c0t6d0                 CD-ROM       connected    configured   unknown
c1                             fc-private   connected    configured   unknown
c1::500000e0106fca91           disk         connected    configured   unknown
c1::500000e0106fcde1           disk         connected    configured   unknown
c1::500000e0106fcf31           disk         connected    configured   unknown
c1::500000e0106fd061           disk         connected    configured   unknown
c1::500000e0106fd7b1           disk         connected    configured   unknown
c1::500000e0106fdaa1           disk         connected    configured   unknown
c1::50800200001d9841           ESI          connected    configured   unknown
c2                             fc-fabric    connected    configured   unknown
c2::201400a0b811804a           disk         connected    configured   unusable
c2::201400a0b8118098           disk         connected    configured   unusable
c2::201700a0b8111580           disk         connected    configured   unusable
c3                             fc-fabric    connected    configured   unknown
c3::201500a0b8118098           disk         connected    configured   unusable
c3::201600a0b8111580           disk         connected    configured   unusable
c3::202500a0b811804a           disk         connected    configured   unusable
c4                             fc-fabric    connected    configured   unknown
c4::200400a0b80f1285           disk         connected    configured   unknown
c4::200400a0b8127a26           disk         connected    configured   unusable
c5                             fc-fabric    connected    configured   unknown
c5::200400a0b82643f5           disk         connected    unconfigured unknown
c5::200500a0b80f1285           disk         connected    configured   unknown
c5::200500a0b8127a26           disk         connected    configured   unusable
c5::200c00a0b812dc5a           disk         connected    configured   unknown
usb0/1                         usb-kbd      connected    configured   ok
usb0/2                         usb-mouse    connected    configured   ok
usb0/3                         unknown      empty        unconfigured ok
usb0/4                         unknown      empty        unconfigured ok
#
```

2. You can also display information about the attachment points on a system. In the following example, c0 represents a fabric-connected host port, and c1 represents a private, loop-connected host port. (Use the **cfgadm** command to manage the device configuration on fabric-connected host ports.)

   By default, the device configuration on private, loop-connected host ports is managed by Solaris host.

   Note: The **cfgadm -1** command displays information about Fibre Channel host ports. You can also use the **cfgadm -al** command to display information about Fibre Channel devices. The lines that include a port World Wide Name (WWN) in the **Ap_ID** field associated with c0 represent a fabric device. Use the **cfgadm configure** and **cfgadm unconfigure** commands to manage those devices and make them available to Solaris hosts.

```
# cfgadm -l
Ap_Id          Type         Receptacle   Occupant     Condition
c0             fc-fabric    connected    unconfigured unknown
c1             fc-private   connected    configured   unknown
```

3. Configure the device using the following command:

   ```
   cfgadm –c configure Ap-Id
   ```

   The **Ap_ID** argument specifies the attachment point ID of the configured Fibre Channel devices. This ID can be the controller number and WWN of a device (for example, `c3::50020f230000591d`).

   See the output example in Step 1 on page 120. Also, see the cfgadm man page for an explanation of attachment points.

   Note: An Ap_Id with type `fc-private` cannot be unconfigured. Only type `fc-fabric` can be configured and unconfigured.

4. Use the **luxadm probe** command to list all mapped LUNs:

```
# luxadm probe
luxadm probe
No Network Array enclosures found in /dev/es

  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006ADE452CBC62d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006ADF452CBC6Ed0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AE0452CBC7Ad0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AE1452CBC88d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AE2452CBC94d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AE3452CBCA0d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AE4452CBCACd0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AE5452CBCB8d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AE6452CBCC4d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AE7452CBCD2d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AE8452CBCDEd0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AE9452CBCEAd0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AEA452CBCF8d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AEB452CBD04d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AEC452CBD10d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006AED452CBD1Ed0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006B2A452CC65Cd0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006B2B452CC666d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006B2C452CC670d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006B2D452CC67Ad0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
       Logical Path:/dev/rdsk/c0t600A0B800011121800006B31452CC6A0d0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c0t600A0B800011121800006B32452CC6ACd0s2
  Node WWN:200400a0b8111218  Device Type:Disk device
    Logical Path:/dev/rdsk/c8t201400A0B8111218d7s2
```

5. You can then use the **luxadm display** *logical path* command to list more details on each mapped LUN, including the number of paths to each LUN. The following example uses a logical path from the previous example.

```
# luxadm display /dev/rdsk/c0t600A0B800011121800006B31452CC6A0d0s2
DEVICE PROPERTIES for disk: /dev/rdsk/c0t600A0B800011121800006B31452CC6A0d0s2
  Vendor:               IBM
  Product ID:           1742-900
  Revision:               0914
  Serial Num:           1T51207691
  Unformatted capacity: 1024.000 MBytes
  Write Cache:          Enabled
  Read Cache:           Enabled
    Minimum prefetch:   0x0
    Maximum prefetch:   0x0
  Device Type:          Disk device
  Path(s):

  /dev/rdsk/c0t600A0B800011121800006B31452CC6A0d0s2
  /devices/scsi_vhci/ssd@g600a0b800011121800006b31452cc6a0:c,raw
   Controller           /devices/pci@7c0/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0
    Device Address              201400a0b8111218,1e
    Host controller port WWN    210100e08ba0fca0
    Class                       secondary
    State                       STANDBY
   Controller           /devices/pci@7c0/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0
    Device Address              201500a0b8111218,1e
    Host controller port WWN    210100e08ba0fca0
    Class                       primary
    State                       ONLINE
   Controller           /devices/pci@7c0/pci@0/pci@8/SUNW,qlc@0/fp@0,0
    Device Address              201400a0b8111218,1e
    Host controller port WWN    210000e08b80fca0
    Class                       secondary
    State                       STANDBY
   Controller           /devices/pci@7c0/pci@0/pci@8/SUNW,qlc@0/fp@0,0
    Device Address              201500a0b8111218,1e
    Host controller port WWN    210000e08b80fca0
    Class                       primary
    State                       ONLINE
#
```

Unconfiguring a failover / failback path

Before you unconfigure a fabric device, stop all activity to the device and unmount any file systems on the fabric device. (See the Solaris administration documentation for unmounting procedures.)

To unconfigure a failover / failback path, complete the following steps:

1. Run the **cfgadm -al** command to display information about the host ports and their attached devices.
2. Unconfigure the LUN by running the following command:

   `cfgadm –c unconfigure` *`Ap-Id`*

   Where *Ap-Id* is the LUN that you want to unconfigure.
3. Run the **cfgadm -al** command again, to make sure that the LUN is now unconfigured.
4. If necessary, define the file structure using the **newfs** command. Add entries to the /etc/vfstab file.

5. Restart the server using the following command:

```
# shutdown -g0 -y -i6
```

Disabling the MPxIO multipath driver

For Solaris 10, unconfigure all devices by typing the `cfgadm –c unconfigure AP-id Ap-id` command. Then, type the `stmsboot –d` command, and accept the default to `Reboot the system now`.

Appendix D. Accessibility

This section provides information about alternative keyboard navigation, which is an accessibility feature of the DS3000 Storage Manager software. 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 alternative 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.

In addition to the keyboard operations that are described in this section, the DS3000 Storage Manager 2 software installation package for Windows includes a screen reader software interface. To enable the screen reader, select **Custom Installation** when you use the installation wizard to install the Storage Manager 2 software on a Windows host or management station. Then, in the Select Product Features window, select **Java Access Bridge** in addition to the other required host software components.

Keyboard focus is not always clearly shown in the panes of the help viewer window. If you cannot see where the keyboard focus is, press Ctrl+F1. If the focus is on the Back, Forward, Print, or Page Setup button in the toolbar, the alternative text for the button is displayed. If alternative text is not displayed, keyboard focus is not on a button. Press Ctrl+Tab to see whether the focus is on one of the navigator tabs (**Contents** tab, **Index** tab, or **Search** tab). If the focus is on one of the navigation tabs, press Shift+Tab to change the focus to the toolbar pane.

Table 10 defines the keyboard operations that you can use 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 component.

In general, navigation among 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 among the components in a group of components

Table 10. DS3000 Storage Manager software alternate keyboard operations

Shortcut	Action
F1	Open the Help.
F10	Move keyboard focus to the main menu bar and posts first menu; use the arrow keys to navigate through the available options.
Alt+F4	Close the management window.
Alt+F6	Move keyboard focus among windows (non-modal) and among management windows.

Table 10. DS3000 Storage Manager software alternate keyboard operations (continued)

Shortcut	Action
Alt+*underlined_letter*	Access menu items, buttons, and other interface components by using the keys that are associated with the underlined letters. For the menu options, press the Alt+*underlined_letter* to access a main menu, and then press the underlined letter to access the individual menu item. For other interface components, press Alt+*underlined_letter*.
Ctrl+F1	Display or conceal a tool tip when keyboard focus is on the toolbar.
Spacebar	Select an item or activate a hyperlink.
End, Page Down	Move keyboard focus to the last item in the list.
Esc	Close the current window (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 among 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.

Appendix E. Getting help and technical assistance

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

Before you call

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

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system and any optional devices are turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system. Information about diagnostic tools is in the *Problem Determination and Service Guide* on the IBM *Documentation* CD that comes with your system.
- Go to the IBM support Web site at http://www.ibm.com/systems/support/ to check for technical information, hints, tips, and new device drivers or to submit a request for information.

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

Using the documentation

Information about your IBM system and preinstalled software, if any, or optional device is available in the documentation that comes with the product. That documentation can include printed documents, 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. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to http://www.ibm.com/systems/support/ and follow the instructions. Also, some documents are available through the IBM Publications Center at http://www.ibm.com/shop/publications/order/.

Getting help and information from the World Wide Web

On the World Wide Web, the IBM Web site has up-to-date information about IBM systems, optional devices, services, and support. The address for IBM System x and xSeries® information is http://www.ibm.com/systems/x/. The address for IBM BladeCenter information is http://www.ibm.com/systems/bladecenter/. The address for IBM IntelliStation® information is http://www.ibm.com/intellistation/.

You can find service information for IBM systems and optional devices at http://www.ibm.com/systems/support/.

Software service and support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with System x and xSeries servers, BladeCenter products, IntelliStation workstations, and appliances. For information about which products are supported by Support Line in your country or region, see http://www.ibm.com/services/sl/products/.

For more information about Support Line and other IBM services, see http://www.ibm.com/services/, or see http://www.ibm.com/planetwide/ for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

Hardware service and support

You can receive hardware service through your IBM reseller or IBM Services. To locate a reseller authorized by IBM to provide warranty service, go to http://www.ibm.com/partnerworld/ and click **Find a Business Partner** on the right side of the page. For IBM support telephone numbers, see http://www.ibm.com/planetwide/. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

IBM Taiwan product service

台灣 IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路 7 號 3 樓
電話：0800-016-888

IBM Taiwan product service contact information:
IBM Taiwan Corporation
3F, No 7, Song Ren Rd.
Taipei, Taiwan
Telephone: 0800-016-888

Appendix F. 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 give you any license to these patents. You can send license inquiries, in writing, to:

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "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 publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml.

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

Intel, Intel Xeon, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc., in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation 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.

Important notes

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity can vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives that are available from IBM.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software might differ from its retail version (if available) and might not include user manuals or all program functionality.

Index

A

B

C

D

M

N

O

P

Q

U

V

W

Z

Part Number: 46M1365

Printed in USA

(1P) P/N: 46M1365