

IBM TotalStorage[®] DS6000



Configuring

IBM TotalStorage[®] DS6000



Configuring

Note:

Before using this information and the product it supports, read the information in "Notices" on page 193.

Twenty-fifth Edition (April 2007)

© Copyright International Business Machines Corporation 2004, 2007. All rights reserved.

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

Contents

Tables	vii
-------------------------	------------

Notices and publication information . . .	ix
--	-----------

Safety notices	ix
Environmental notices	ix
Product recycling and disposal	ix
Battery return program	x
How to send your comments	xi

Chapter 1. Configuring	1
---	----------

Chapter 2. Creating an express configuration	3
---	----------

Creating open systems volumes using express configuration	3
Creating zSeries volumes using express configuration	4
Creating iSeries volumes using express configuration	5

Chapter 3. Creating a custom logical storage configuration	9
---	----------

Creating arrays	10
Creating ranks	11
Creating extent pools	11
Creating zSeries LCUs.	13
Creating zSeries volumes.	13
Creating open systems volumes	14
Creating open systems volume groups	16
Applying a configuration (simulated only)	17
Enabling host system I/O	18

Chapter 4. Managing logical storage configurations	19
---	-----------

Viewing status (real-time only)	19
Viewing and modifying properties.	19
Viewing array disk drive module information	19
Adding arrays to ranks	19
Modifying ranks.	20
Repairing a rank (real-time only)	20
Adding ranks to extent pools	20
Removing ranks from extent pools	21
Modifying extent pools	21
Modifying open systems volume groups.	21
Modifying zSeries volumes	22
Modifying LCUs.	23

Chapter 5. Deleting a logical storage configuration	25
--	-----------

Deleting host systems	25
Deleting open systems volume groups	25
Deleting arrays	26
Deleting ranks	26
Deleting open systems volumes	27
Deleting extent pools	27

Deleting zSeries volumes	28
Deleting logical control units (LCUs)	28

Chapter 6. DS6000 host attachment . . .	29
--	-----------

General requirements for attaching a host system.	29
Downloading and installing a host adapter driver	30
Host attachment general considerations	33
IBM System Storage Multipath Subsystem Device Driver	33
Attachment restrictions for Copy Services	33
DS6000 fabric zoning	34
Apple Macintosh host attachment	35
Fujitsu PRIMEPOWER host attachment	35
Locating the WWPN for a Fujitsu PRIMEPOWER host	36
Installing the Emulex adapter for a PRIMEPOWER host	36
Hewlett-Packard Server (HP-UX) host attachment	40
Locating the WWPN for an HP-UX host.	40
Setting the queue depth for an HP-UX host.	41
Configuring clustering on an HP-UX host	41
HP AlphaServer or Integrity server OpenVMS host attachment	43
Confirming the installation of the OpenVMS operating system	43
Installing the KGPSA-xx adapter card in an AlphaServer OpenVMS host.	44
Setting the mode for the KGPSA-xx host adapter in an AlphaServer OpenVMS host.	45
Locating the WWPN for an OpenVMS host.	45
Managing devices for an Integrity server using the OpenVMS EFI	47
Confirming fibre-channel switch connectivity for OpenVMS	48
Defining a UDID for OpenVMS storage unit volumes	49
Accessing the storage unit volumes from an OpenVMS host	51
Guidelines for OpenVMS LUN 0	52
Restrictions for OpenVMS fibre-channel connections	53
HP AlphaServer Tru64 UNIX host attachment	54
Confirming the installation of the Tru64 UNIX operating system	55
Installing the KGPSA-xx adapter card in a Tru64 UNIX host system	55
Setting the mode for the KGPSA-xx host adapter in a Tru64 UNIX host system	56
Locating the WWPN for a Tru64 UNIX host	56
Configuring a Tru64 UNIX host	57
Configuring storage for a fibre-channel Tru64 UNIX host.	62
Removing persistent reserves for Tru64 UNIX 5.x	62
Limitations for Tru64 UNIX	65
IBM eServer iSeries host attachment	66

Locating the WWPN for an IBM eServer iSeries host	67	Configuring an Intel or AMD host running VMware ESX Server	133
LUN considerations for IBM iSeries hosts	67	Copy Services considerations for VMware ESX servers	136
Switch support for IBM iSeries hosts	68	Intel or AMD running Microsoft Windows 2000 or 2003 host attachment	137
Recommended configurations for IBM iSeries hosts	68	Installing an Emulex adapter and driver for Windows 2000 or 2003	137
Running the Linux operating system on an IBM i5 server	70	Installing a Netfinity adapter and driver for Windows 2000 or 2003	141
IBM eServer pSeries or IBM RS/6000 host attachment	74	Installing a QLogic adapter and driver for Windows 2000 or Windows 2003	143
Installing the host attachment package on IBM pSeries AIX hosts from the CD	75	Updating the Windows 2000 or 2003 device driver	145
Installing the host attachment package on IBM pSeries AIX hosts from the SDD Web site	76	Locating the WWPN for a Windows 2000 or 2003 host	145
Locating the WWPN for an IBM eServer pSeries or an RS/6000 host	77	Verifying that Windows 2000 or 2003 is configured for storage	146
Attaching a pSeries or RS/6000 host	77	Setting the TimeOutValue registry for Windows 2000 or 2003	146
Support for fibre-channel boot on a pSeries host	78	Installing remote fibre-channel boot support for a Windows 2000 or 2003 host system	147
Attaching to multiple RS/6000 or pSeries hosts without a HACMP host	80	Intel running Novell NetWare host attachment	149
Installing the subsystem device driver on an IBM pSeries host running Linux	81	Installing an Emulex adapter and driver for a Novell NetWare host	149
IBM eServer zSeries or IBM S/390 host attachment	85	Installing a QLogic adapter and driver for a Novell NetWare host	150
FICON-attached S/390 and zSeries hosts overview	86	Silicon Graphics host attachment	152
Linux for S/390 and zSeries	88	Checking the version of the IRIX operating system	153
Registered state-change notifications (RSCNs) on zSeries hosts	94	Installing a fibre-channel adapter card for the SGI host system	153
Analyzing service information messages for S/390 and zSeries systems	94	Installing an optical cable for SGI in a switched-fabric topology	154
IBM NAS Gateway 500 host attachment	95	Installing an optical cable for SGI in an arbitrated loop topology.	154
Locating the WWPN on a NAS Gateway 500 host system	95	Confirming switch connectivity for SGI.	154
IBM SAN File System host attachment	97	Displaying zoning information for the switch	155
IBM SAN Volume Controller host attachment	98	Locating the WWPN for an SGI host	156
Intel or AMD running Linux host attachment	98	Confirming storage connectivity on a SGI host	156
Installing an Emulex adapter on an Intel or AMD host running Linux	99	Configuring the storage unit for host failover using IRIX multipathing.	158
Installing a QLogic adapter on an Intel or AMD host system running Linux	103	Configuring storage in a switched fabric or arbitrated loop topology.	160
Locating the WWPN for a Linux host	106	Sun host attachment	164
Managing SCSI disk connectivity.	107	Installing fibre-channel adapters and adapter drivers on a Sun host.	165
Creating a DS6000 file system and partitions for an Intel host running Linux	116	Configuring host device drivers on a Sun host	167
Assigning the system ID to the partition of an Intel or AMD host running Linux	117	Setting system parameters on a Sun host	183
Creating a file system for an Intel or AMD host running Linux	118	Locating the WWPN for a Sun host	185
Fibre-channel boot on an Intel or AMD host running Linux	119	Attaching a Sun host using Storage Traffic Manager System	186
Intel or AMD running VMware ESX server host attachment	127	Attaching a Sun host using Sun Cluster	187
Installing the Emulex adapter for an Intel or AMD host running VMware ESX Server	127	iSCSI Gateway host attachment	188
Installing the QLogic adapter and driver for an Intel or AMD host running VMware ESX Server.	128	Attachment overview of the iSCSI Gateway host	189
Defining the number of disk devices on a VMware ESX server	129	Ethernet adapter attachment considerations for the iSCSI Gateway host	189
SCSI disk considerations for an Intel or AMD host running VMware ESX server	129	Configuring for storage for the iSCSI Gateway host	190
		iSCSI Gateway operation through the IP Service Module	190

Chapter 7. Network configurations 191

Notices 193

Accessibility	194
Trademarks	195
Terms and conditions.	196
Electronic emission notices	197
Federal Communications Commission (FCC) statement.	197
Industry Canada compliance statement.	197

European Union EMC Directive conformance statement.	197
Japanese Voluntary Control Council for Interference (VCCI) class A statement	199
Korean Ministry of Information and Communication (MIC) statement.	199
Taiwan class A compliance statement	199

Index 201

Tables

1.	Host adapter driver download steps by vendor	30	12.	Recommended settings for the QLogic QLA23xx, and QLA246x adapters for a Novell NetWare host	150
2.	Examples of SDD MPIIO path control modules	33	13.	SGI qlfc_use_connection_mode parameters	154
3.	Recommended configuration file parameters for the Emulex LP9002L adapter	37	14.	Recommended configuration settings for AMCC adapters on a Sun host	170
4.	Recommended Tru64 UNIX settings for the SCSI DEVICE..	60	15.	Recommended configuration settings for Emulex adapters on a Sun host	172
5.	Capacity and models of disk volumes for IBM iSeries	69	16.	Recommended configuration settings for Sun-Emulex adapters on a Sun host	174
6.	Recommended settings for the QLogic adapter card for an Intel or AMD host running Linux	104	17.	System log parameters for configuring Sun-Emulex adapter messages	176
7.	Recommended settings for the QLogic adapter for an Intel or AMD host running VMware ESX Server	128	18.	Recommended configuration settings for QLogic QLA23xx adapters with driver level 4.03 on a Sun host	177
8.	Recommended configuration file parameters for the Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, LP10000DC, LP11000, LP11002, LPE11000, and LPE11002 adapters	138	19.	Recommended configuration settings for QLogic QLA23xx and QLogic QLA246x adapters with a San Surf configuration on a Sun host	178
9.	StorPort Miniport driver settings	141	20.	Recommended configuration settings for Sun-QLogic adapters on a Sun host	180
10.	Recommended settings for the QLogic QLA23xx, QLA246x, and QLE246x adapters for Windows 2000 or Windows 2003	143	21.	Recommended Sun settings for the SCSI device	187
11.	Jumper change for QLA2340/2342, and an IBM 2005-B32 (Brocade Silkworm 4100) running Windows 2003	145			

Notices and publication information

This section contains information about safety notices that are used in this guide, environmental notices for this product, publication information, and information about sending your comments to IBM.

Safety notices

Complete this task to find information about safety notices.

To find the translated text for a danger or caution notice:

1. Look for the identification number at the end of each danger notice or each caution notice. In the following examples, the numbers **1000** and **1001** are the identification numbers.

DANGER

A danger notice indicates the presence of a hazard that has the potential of causing death or serious personal injury.

1000

CAUTION:

A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury.

1001

2. Find the number that matches in the *IBM System Storage Solutions Safety Notices for IBM Versatile Storage Server and IBM System Storage Enterprise Storage Server, GC26-7229*.


Environmental notices

This section identifies the environmental guidelines that pertain to this product.

Product recycling and disposal

This unit contains recyclable materials.

This unit must be recycled or discarded according to applicable local and national regulations. IBM® encourages owners of information technology (IT) equipment to responsibly recycle their equipment when it is no longer needed. IBM offers a variety of product return programs and services in several countries to assist equipment owners in recycling their IT products. Information on IBM product recycling offerings can be found on IBM's Internet site at <http://www.ibm.com/ibm/environment/products/prp.shtml>.



Notice: This mark applies only to countries within the European Union (EU) and Norway.

Appliances are labeled in accordance with European Directive 2002/96/EC concerning waste electrical and electronic equipment (WEEE). The Directive determines the framework for the return and recycling of used appliances as applicable throughout the European Union. This label is applied to various products to indicate that the product is not to be thrown away, but rather reclaimed upon end of life per this Directive.


In accordance with the European WEEE Directive, electrical and electronic equipment (EEE) is to be collected separately and to be reused, recycled, or recovered at end of life. Users of EEE with the WEEE marking per Annex IV of the WEEE Directive, as shown above, must not dispose of end of life EEE as unsorted municipal waste, but use the collection framework available to customers for the return, recycling and recovery of WEEE. Customer participation is important to minimize any potential effects of EEE on the environment and human health due to the potential presence of hazardous substances in EEE. For proper collection and treatment, contact your local IBM representative.

Battery return program

This product may contain sealed lead acid, nickel cadmium, nickel metal hydride, lithium, or lithium ion battery. Consult your user manual or service manual for specific battery information. The battery must be recycled or disposed of properly. Recycling facilities may not be available in your area. For information on disposal of batteries outside the United States, go to <http://www.ibm.com/ibm/environment/products/index.shtml> or contact your local waste disposal facility.

In the United States, IBM has established a return process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and other battery packs from IBM Equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426-4333. Please have the IBM part number listed on the battery available prior to your call.

In the Netherlands the following applies:



For Taiwan:



Please recycle batteries.

廢電池請回收

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this information or any other DS6000™ series documentation, you can submit them in the following ways:

- e-mail

Submit your comments electronically to the following e-mail address:

starpubs@us.ibm.com

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

- Mail

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

International Business Machines Corporation
RCF Processing Department
Department 61C
9032 South Rita Road
TUCSON AZ 85775-4401

Chapter 1. Configuring

The topics in this section provide information related to configuring your DS6000. Topics covered include express and custom logical storage configuration as well as network configurations.

Chapter 2. Creating an express configuration

This section contains the high-level roadmaps for creating express logical configurations.

You can create the following express logical storage configurations:

- Open systems logical storage configurations
- zSeries logical storage configurations
- iSeries logical storage configurations

Creating open systems volumes using express configuration

Complete this task to quickly configure a storage complex with open systems volumes.

Ensure that you have obtained and applied an Operating Environment License with a maximum storage amount that exceeds the amount of storage that you have configured in your simulated configuration.

To create open systems volumes using express configuration, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Express configuration wizard**. The Express configuration wizard page appears.
2. From the **Select storage unit** list, select the storage unit for the volumes that you are configuring.
3. From the **Select volume type** list, select **Standard Open System (FB)**, and then click **Next**. The Open systems volumes page appears.
4. Select either **RAID 5** or **RAID 10** under **Select RAID type**.
5. Make selections for two of the following three fields and then click **Calculate** configure the amount of storage:
 - Select a value from the **Amount of available storage to use** drop-down list. If you select **User defined GB**, you must enter a value, in GB, in the **User defined GB** field. If you select **User defined %**, you must enter a percentage value in the **User defined %** field.
 - Enter an amount, in GB, in the **Volume size** field to determine the size of the individual volumes. The DS Storage Manager automatically determines the number of volumes that are created.
 - Enter an amount, as an integer, in the **Volume quantity** field to determine the number of individual volumes. The DS Storage Manager automatically determines the size of the volumes that are created.
6. To enable host creation for the open system volumes, select **Create host**.
7. Click **Next**. The Set volume naming page is displayed.
8. To specify a sequence of volume names for the quantity that you created, select **Generate a sequence of nicknames based on the following**. Select the **Use hexadecimal sequence** check box to generate a hexadecimal sequence. See Using hexadecimal values for conversion information.
9. Select **Verify nicknames are unique** to determine if volume names that are generated on this page are unique.

10. Enter the prefix (alphabetic) that you want for the volumes in the **Prefix** field.
11. Enter the suffix (numeric) that you want for the volumes in the **Numeric suffix** field. Click **Next**. If you selected **Create host**, the Group volumes page appears.
12. Enter a name (up to 16 characters) for the volume group in the **Volume group name** field, select the volumes to include in the volume group under **Select volumes**, and then click **Next**.
 - If you did not select to create a host, the Verification page appears and you can move to the final step.
 - If you selected to create a host, the General host information page appears and you can continue to the next step.
13. Select the host system for the open systems volumes in the **Host system** field, and then enter a nickname (up to 16 characters) for the host in the **Nickname** field.
14. (Optional) Enter a description (up to 256 characters) for the host, and then click **Next**. The Host ports page appears.
15. Enter the number of host ports to configure for the open systems volumes in the **Quantity** field, and then select the host attachment type from the **Type** list.
16. To group the ports that you are configuring and manage them as a single host attachment, select **Group ports to share a common set of volumes**, and then click **Next**. The Define WWPNS page appears.
17. Select or enter the 16-digit WWPNS for each host port that you want to include in the identifier, and then click **Next**. The Assign host to volume group page is displayed.
18. To map the volume group to the host attachment, select **Assign host attachment to volume group**, and then click **Next**.
19. Review the details of the configuration on the Verification page. You can navigate through the Express Configuration pages to make any changes to the configuration by clicking **Back** or **Next**, or by selecting a specific step in the wizard in the left navigation. When you are satisfied with the details of the configuration, click **Finish**.

Creating zSeries volumes using express configuration

Complete this task to quickly and easily configure a storage complex with zSeries volumes.

Ensure that you have obtained and applied an Operating Environment License with a maximum storage amount that exceeds the amount of storage that you have configured in your simulated configuration.

Note: Before you configure, disable any firewalls, as they might interfere with DS6000 communication.

To create zSeries volumes using express configuration, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Express configuration wizard**. The Express configuration wizard page appears.
2. From the **Select storage unit** list, select the storage unit for the volumes that you are configuring.

3. In the **Select volume type** list, select **zSeries (CKD)**, and then click **Next**. The **zSeries volumes** page appears.
4. Select either **RAID 5** or **RAID 10** under **Select RAID type**.
5. Select the appropriate model in the **Volume type** list.
6. Select one of the following options to configure an amount of available storage:
 - Select **Amount of available storage to use**. Either select a percentage from the list or enter a value in one of the **User defined** fields representing the space to configure.
 - Select **Volume quantity**, and enter the number of volumes that you want to create.
7. Click **Calculate**. The value for the item that you did not select is automatically displayed.
8. Click **Next**. The **LCU/SSID** page appears.
9. Select the LCU that you want as the starting LCU in the **Select starting LCU** list.
10. Select the type for the LCU in the **LCU type** list.
11. If you want to change the default SSID, highlight the SSID in the **SSID** field, and type your selection.
12. If the **PAVs** section appears, optionally select **Define number of aliases per base** and enter a number in the **Aliases/base** field.
13. Enter the appropriate number in the **Number of LCUs to create** field, and then select one of the following options:
 - Select **Spread volumes equally across LCUs**.
 - Select **Utilize all addresses in each LCU** to have the quantity of base volumes that are created use all the addresses in the specified number of LCUs.
14. Click **Next**. The **Set volume naming** page appears.
15. To specify a sequence of volume names for the quantity created, select **Generate a sequence of nicknames based on the following**.
16. Select **Verify nicknames are unique** to determine whether volume names that are generated on this page are unique.
17. Enter the prefix (alphabetic) that you want for the volumes in the **Prefix** field.
18. Enter the suffix (numeric) that you want for the volumes in the **Numeric suffix** field. If you are using hexadecimal values, see *Using hexadecimal values for conversion information*. Click **Next**.
19. Review the details of the configuration on the **Verification** page. You can navigate through the **Express Configuration** pages to make any changes to the configuration by clicking **Back** or **Next**, or by selecting a specific step in the wizard in the left navigation. When you are satisfied with the details of the configuration, click **Finish**.

Creating iSeries volumes using express configuration

Complete this task to quickly configure a storage complex with iSeries volumes.

Ensure that you have obtained and applied an Operating Environment License with a maximum storage amount that exceeds the amount of storage that you have configured in your simulated configuration.

Note: Before you configure, disable any firewalls, as they might interfere with DS6000 communication.

To create iSeries volumes using express configuration, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Express configuration wizard**. The Express configuration wizard page appears.
2. From the **Select storage unit** list, select the storage unit for the volumes that you are configuring.
3. In the **Select volume type** list, select **iSeries (FB)**, and then click **Next**. The iSeries volumes page appears.
4. Select either **RAID 5** or **RAID 10** under **Select RAID type**.
5. Select either **Protected** or **Unprotected** under **Select volume type**.
6. Select one of the provided iSeries volume sizes from the **Select volume size** field.
7. Select one of the following options to configure an amount of available storage:
 - Select **Amount of available storage to use** and either select a percentage from the list or enter a value in one of the **User defined** fields representing the space to configure.
 - Select **Volume quantity** and enter the number of volumes that you want to create.
8. Click **Calculate**. The value for the item that you did not select is automatically calculated.
9. To enable host creation for the iSeries volumes, select **Create host**.
10. Click **Next**. The Set volume naming page is displayed.
11. To specify a sequence of volume names for the quantity created, select **Generate a sequence of nicknames based on the following**.
12. Select **Verify nicknames are unique** to determine if the volume names that are generated on this page are unique.
13. Enter the prefix (alphabetic) that you want for the volumes in the **Prefix** box.
14. Enter the suffix (numeric) that you want for the volumes in the **Numeric suffix** box. If you are using hexadecimal values, see Using hexadecimal values for conversion information. Click **Next**.
15. Enter a name (up to 16 characters) for the volume group in the **Volume group name** box, select the volumes to include in the volume group under **Select volumes** and click **Next**. The General host information page is displayed.
 - If you did not select to create a host, the Verification page is displayed and you can move to the final step.
 - If you selected to create a host, the General host information page is displayed and you can continue to the next step.
16. Select the host system for the iSeries volumes in the **Host system** box, and then enter a nickname (up to 16 characters) for the host in the **Nickname** box.
17. (Optional) Enter a description (up to 256 characters) for the host and click **Next**. The Host ports page is displayed.
18. Enter the number of host ports to configure for the iSeries volumes in the **Quantity** box, and then select the host attachment type from the **Type** list.
19. To group the ports that you are configuring and to manage them as a single host attachment, select **Group ports to share common set of volumes** and click **Next**. The Define WWPNs page is displayed.

20. Select or enter the 16-digit WWPNs for each host port that you want to include in the identifier and click **Next**. The Assign host to volume group page is displayed.
21. To map the volume group to the host attachment, select **Assign host attachment to volume group** and click **Next**.
22. Review the details of the configuration on the Verification page. You can navigate through the Express Configuration pages to make any changes to the configuration by clicking **Back** or **Next**, or by selecting a specific step in the wizard in the left navigation. When you are satisfied with the details of the configuration, click **Finish**.

Chapter 3. Creating a custom logical storage configuration

Complete this task to create a simulated or real-time custom logical configuration where you manually determine the storage allocation details. You must complete all of the steps and substeps in this task to create your configuration.

You must have installed and configured the DS Storage Manager before you can create your custom logical configuration. You also must have obtained and applied an Operating Environment License (OEL) with a maximum storage amount that exceeds the amount of storage that you intend to configure. In addition, your network must be configured so that it can support all the components and functions that you will use with your storage unit.

If you are working in simulated mode, you must create a new storage configuration or open an existing storage configuration before completing this task. Follow the steps in Creating a configuration file (simulated only) to manage your configuration files.

You can create either a simulated or real-time custom logical configuration by following the steps in this task. A real-time configuration allows you to manage physical and logical configurations from existing storage complexes, and storage servers over your network in real time. A simulated configuration allows you to create or import a new *simulated* instance of your DS6000 series.

Note: Before you configure, disable any firewalls, as they might interfere with DS6000 communication

Return to the next step of this task after completing each of the subtasks below.

1. Create your arrays. To create an array, you must complete the subtask in "Creating arrays" on page 10.
2. Create your ranks. To create a rank, you must complete the subtask in "Creating ranks" on page 11.
3. Create your extent pools. To create an extent pool, you must complete the subtask in "Creating extent pools" on page 11.
4. Perform one of the following tasks for either your zSeries or open systems host:
 - If you are working with a zSeries host, you must create zSeries LCUs and zSeries volumes.
 - a. Create zSeries LCUs. To create zSeries LCUs, you must complete the subtask in "Creating zSeries LCUs" on page 13.
 - b. Create zSeries volumes. To create zSeries volumes, you must complete the subtask in "Creating zSeries volumes" on page 13.
 - If you are working with an open systems host, you must create open systems volumes and open systems volume groups.
 - a. Create open systems volumes. To create open systems volumes, you must complete the subtask in "Creating open systems volumes" on page 14.
 - b. Create open systems volume groups. To create open systems volume groups, you must complete the subtask in "Creating open systems volume groups" on page 16.

5. Reconfigure your host system to start sending I/O to the DS6000. To reconfigure your host system, you must complete the subtask in “Enabling host system I/O” on page 18.

Note: If you are working in simulated mode, you must apply the configuration before you can complete this step. To apply the configuration, complete the subtask in “Applying a configuration (simulated only)” on page 17.

Creating arrays

Complete this task to create arrays, either automatic or custom, and to specify their RAID and rank attributes.

Perform the following steps to create arrays:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Arrays**. On the Arrays main page, in the **Select storage unit** drop-down list, select a storage unit.
 2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create Array – Definition method page is displayed.
 3. In Create Array – Definition method, choose to either create your arrays automatically, where the DS Storage Manager specifies the array sites, or to create custom arrays where you specify the array sites. Click **Next**.
 - If you chose to create your arrays automatically, in Create Array – Array configuration (Auto), specify the quantity and RAID type of the arrays that you are creating. You can optionally select **Create an 8 disk array** to create an array that contains two, four-disk array sites with the same specifications. Click **Next** after you have made your selections. The Create Array – Add arrays to ranks page is displayed.
 - If you chose to create custom arrays, perform the following steps:
 - a. In Create Array – Array configuration (Custom), specify the RAID type for the arrays that you are creating and select the array site numbers that you want to assign to the arrays. You can optionally select **Create an 8 disk array** to create arrays that contain two, four-disk array sites with the same specifications. Click **Next** after you have made your selections.
 - b. If you selected to create an 8 disk array, the Create Array: Real-time – Second array-site selection page is displayed. Choose the second array site for each array that was selected in the Create Array – Array configuration (Custom) page. Secondary array sites cannot be shared among arrays. Click **Next**. The Create Array – Add arrays to ranks page is displayed.
- Note:** When configuring two 8 disk arrays (one Raid 5, one Raid 10), you must create the Raid 10 array first and make sure that it is being created on the array sites that contain spares. Then create the Raid 5.
4. In Create Array – Add arrays to ranks, you can optionally choose to add the arrays that you are creating to ranks. If you choose to add your arrays to ranks, you must specify the storage type, FB or CKD, for the ranks that you are creating. Click **Next**. The Create Array – Verification page is displayed.
 5. In Create Array – Verification, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and then specify the correct values. Otherwise, click **Finish** to complete the array creation process.

If you are creating a new logical storage configuration, return to Chapter 3, “Creating a custom logical storage configuration,” on page 9 and complete the rest of the steps.

Creating ranks

Complete this task to create ranks and to specify their arrays and extent pools.

Note: If you are building a configuration in simulated mode or you are importing a real-time configuration into the simulated configurator, extent allocation per rank in an extent pool might not be the same. This means that the rank extent allocations that are created in simulated mode or imported as part of a real-time configuration might not match the allocation when you compare applied and imported configurations to real-time configurations at the rank level, although the extent usage is equivalent. The real-time configuration extent allocation at the rank level will be configured appropriately.

Perform the following steps to create ranks:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Ranks**. On the Ranks main page, in the **Select storage unit** drop-down list, select a storage unit.
2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create Rank – Select array for rank page is displayed.
3. In the **Select** column, select an array. Then click **Next**. The Create Rank – Define rank properties page is displayed. The **Rank number** is provided by default.

Note: If there are no available arrays for selection, you can create one by clicking the **Create new array** button. When you have finished creating an array, you are returned to this page, and the array that you created is available for selection.

4. Specify the **Storage type** and then click **Next**. The Create Rank – Select extent pool page is displayed.
5. Optionally, you can select one extent pool. In the **Select** column, select an extent pool. Then click **Next**. The Create Rank – Verification page is displayed.

Note: If there are no available extent pools for selection, you can create one by clicking the **Create new extent pool** button. When you have finished creating an extent pool, you are returned to this page, and the extent pool that you created is available for selection.

6. In Create Rank – Verification, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and then specify the correct values. Otherwise, click **Finish** to complete the rank creation process.

If you are creating a new logical storage configuration, return to Chapter 3, “Creating a custom logical storage configuration,” on page 9 and complete the rest of the steps.

Creating extent pools

Complete this task to create extent pools, automatic or custom, and to specify extent pool parameters.

Each extent pool is either CKD or FB. An extent pool is associated with either server 0 or server 1. You need a minimum of 2 extent pools, one for each server, to make full use of the resources. An extent pool is divided into equal extents of the following size:

- FB - 1GB
- CKD - 1113 cylinders

Your maximum number of extent pools cannot exceed the number of ranks. It is recommended that you put only ranks with the same capacity and speed into the same extent pool. RAID types cannot be intermixed in an extent pool. For performance management, you can create extent pools of a single rank.

Perform the following steps to create extent pools:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Extent pools**. On the Extent pools main page, in the **Select storage unit** drop-down list, select a storage unit.
2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create Extent Pool – Define properties page is displayed.
3. Choose either to create your extent pool automatically based on storage requirements (where the DS Storage Manager automatically selects your arrays and ranks) or to create a custom extent pool, where you decide which arrays and ranks are assigned to the extent pools.

- If you chose to create your extent pools automatically, in Create Extent Pool – Define requirements, provide a nickname for your extent pool, select the storage type (either FB or CKD) and RAID type, and enter the required amount of storage. You can optionally select to use any existing unassigned arrays and ranks if you want the DS Storage Manager to use arrays and ranks that already exist. Click **Next**. The Create Extent Pool – Reserve storage page is displayed.

Note: An extent pool uses all of the available storage in arrays or ranks that make up the extent pool, even if you specify a storage amount that is less than the entire array or rank storage amount. For example, if you have a 385 GB rank and need 100 GB of storage, your extent pool will contain 385 GB of storage. However, if you have two ranks with 100 GB of storage in each rank and you need 120 GB in your extent pool, your extent pool will contain two ranks with 200 GB of storage.

- If you chose to create a custom extent pool, perform the following steps:
 - a. In Create Extent Pool – Define properties, provide a nickname for your extent pool, then select the storage type (either FB or CKD), RAID type, and server that are assigned to the extent pool. Click **Next** after you have made your selections.
 - b. In Create Extent Pool – Select ranks, select the rank that you want to assign to the extent pool or create a new rank to assign to the extent pool. Click **Next**. The Create Extent Pool – Reserve storage page is displayed.
- 4. In the Create Extent Pool – Reserve storage page, optionally enter the amount of storage that you want the DS Storage Manager to reserve for the extent pool. Click **Next**. The Create Extent Pool – Verification page is displayed.
- 5. In the Create Extent Pool – Verification page, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back**, as appropriate, to return and to specify the correct values. Otherwise, click **Finish** to complete the extent pool creation process.

If you are creating a new logical storage configuration, return to Chapter 3, “Creating a custom logical storage configuration,” on page 9 and complete the rest of the steps.

Creating zSeries LCUs

Complete this task to create zSeries logical control units (LCUs) and to specify their attributes and properties.

To create zSeries LCUs, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **zSeries** → **LCUs**. On the LCUs main page, from the **Select storage unit** list, select a storage unit.
2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create LCU – Select from available LCUs page is displayed.
3. Select one or more LCUs from the list of those available. Click **Next** to continue. The Define LCU properties page is displayed.

Note: If you are creating LCUs to associate with a specific extent pool, you must select LCU ID numbers that match, as even or odd, the server number that was specified for the extent pool when it was created. For example, an extent pool that is associated with Server 1 can be associated only with LCUs that have odd numbers.

4. In the Define LCU properties page, define the parameters for the selected LCUs.
 - a. You can change the default SSID by highlighting and typing over it.
 - b. Specify the LCU types and establish the timeout times in seconds.
5. Click **Next** to continue. The Verification page is displayed.
6. Use the Verification page to review the established attributes and verify that they are correct. If the attributes and values are not correct, click **Back**, as appropriate, to return and to specify the correct values. Otherwise, click **Finish** to complete the LCU creation process.

If you are creating a new logical storage configuration, return to Chapter 3, “Creating a custom logical storage configuration,” on page 9 and complete the rest of the steps.

Creating zSeries volumes

Complete this task to create zSeries volumes and to specify their attributes and properties.

Only certain LCUs are available for an extent pool. To be available, the LCU numbers must match in parity—even or odd—with the server number that is associated with the extent pool. For example, given an extent pool that was associated with Server 1, which is odd, you will have LCUs for selection in the Define volume characteristics page only if there are LCUs that are also odd. If none appear, you must create LCUs by assigning odd numbers to them in the Create zSeries LCUs wizard.

To create zSeries volumes, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **zSeries** → **Volumes – zSeries**. On the Volumes – zSeries

main page, from the **Select storage unit** list, select a storage unit. In the **Select LCU** drop-down list, select a zSeries LCU.

2. From the **Select Action** list, select **Create...**, and then click **Go**. The Create zSeries volume – Select extent pool is displayed.
3. Select the extent pools that you want to use to create your volumes or click **Create new extent pool** to create new CKD extent pools that you can then use to create your zSeries volumes. Click **Next**. The Create zSeries volume – Define base volume characteristics page is displayed.
4. In Create zSeries volume – Define base volume characteristics, specify the volume type and LCU number that are associated with the zSeries volumes. Click **Next**. The Create zSeries volume – Define base volume properties page is displayed.

Note: LCUs are displayed for selection only if their numbers match in parity, even or odd, with the server number that was selected for the extent pool. If no LCUs are displayed, you must create them with the correct parity before you can continue.

5. In Create zSeries volume – Define base volume properties, select an addressing policy and if necessary, supply a maximum LCU address. Enter the quantity of base volumes. If you selected a custom volume type on the Define base volume characteristics page, you have the option to specify a size in cylinders. You also must specify the quantity of LCUs, the base start address, and a sorting order of Ascending or Descending. The Available storage table shows the amount of storage that you can work with while defining your volume properties. Click **Next**. The Create zSeries volume – Create volume nicknames page is displayed.
6. In Create zSeries volume – Create volume nicknames, choose if you want to generate a sequence of nicknames based on the specified prefix and suffix and if you want to use hexadecimal values. If you are using hexadecimal values, see Using hexadecimal values for conversion information. Click **Next**. The Create zSeries volume – Define alias assignments page is displayed.
7. In Create zSeries volume – Define alias assignments, select the LCUs for which you want to define the aliases. Then specify the starting address and address sorting order. Provide the values for the number of aliases per number of volumes. For example, if you enter 1 in the **Aliases** field, and 4 in the **Per volumes field**, you are assigning one alias name for every four volumes. Click **Next** to continue. The Create zSeries volume – Verification page is displayed.
8. Use the Verification page to review the attributes and verify that they are correct. If the attributes and values are not correct, click **Back**, as appropriate, to return and specify the correct values. Otherwise, click **Finish** to complete the volume creation process.

If you are creating a new logical storage configuration, return to Chapter 3, “Creating a custom logical storage configuration,” on page 9 and complete the rest of the steps.

Creating open systems volumes

Complete this task to create open systems volumes and to specify their attributes and properties.

Note: Before you configure, disable any firewalls, as they might interfere with DS6000 communication.

To create open systems volume groups, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Open systems** → **Volumes – Open systems**. On the Volumes – Open systems main page, in the **Select storage unit** list, select a storage unit.
2. Select a volume. From the **Select Action** list, select **Create...**, and then click **Go**. The Create Volume – Select extent pool page is displayed.
3. You must select an extent pool for the target volumes. After you select the extent pool, click **Next**. The Create Volume – Define volume characteristics page is displayed.

Note: If there are no available extent pools available for selection, you can create one by clicking the **Create new extent pool** button. When you have finished creating an extent pool, you are returned to this page, and the extent pool that you created is available for selection.

4. In Create Volume – Define volume characteristics, define the volume type for the new volumes. You can select any number of volume groups from the **Select volume groups** list to associate with the new volumes. After you define the volume characteristics, click **Next**. The Create Volume – Define volume properties page is displayed.
5. Define the volume properties. To calculate the maximum size, enter a value in the **Quantity** field and click the **Calculate max size** button. To calculate the maximum quantity, enter a value in the **Size** field and click the **Calculate max quantity** button. The **Extent size** and **Available Storage Capacity in Extent Pool** fields are displayed in the units that correspond to the selected volume type.

If you click the **Calculate max size** button, the **Size** field is populated with the calculated value. You can overwrite the value and enter an integer value for a lesser size. If you selected one of the iSeries volume types on the previous Define volume characteristics page and Decimal GB (10^9 bytes) for the capacity units, the values here include 8.56, 17.54, 35.16, 36.00, 70.56, 141.12, and 282.25. If you selected one of the iSeries volume types on the previous Define volume characteristics page and Binary GB (2^{30} bytes) for the capacity units, the values here include 8.00, 16.34, 32.75, 33.53, 65.72, 131.44, and 243.80.

If you select the **Select LSSs for volumes** checkbox, you are required to select from the list of available LSSs for these open systems volumes.

After you define the volume properties, click **Next**. The Create Volume – Create volume nicknames page is displayed.

6. In Create Volume – Create volume nicknames, you can optionally create one or more nicknames that are based on entries in the **Prefix** or **Suffix** fields. If you do not specify nicknames, only a volume number is created. Click **Next**. The Create Volume – Verification page is displayed.

Tip: If you plan to create volume groups, you can use a unique and meaningful nickname that can help you easily find the volumes that you want to include in a volume group. For example, if you are creating multiple volumes of data for a specific department, you can make the nickname prefix an abbreviation of that department's name. Then you can use a predetermined range of numbers in the nickname suffix to identify each individual volume.

7. In Create Volume – Verification, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and to specify the correct values. Otherwise, click **Finish** to complete the volume creation process.

If you are creating a new logical storage configuration, return to Chapter 3, “Creating a custom logical storage configuration,” on page 9 and complete the rest of the steps.

Creating open systems volume groups

Complete this task to create open systems volume groups and to specify their properties and parameters.

Volume groups provide a volume masking function. When you create a volume, you can assign it to one or more volume groups. Open systems server attachments are assigned to volume groups to control access to the volumes. There can be up to 8320 volume groups. Any host port (or host port group) can access only one volume group.. A volume group can be accessed by multiple servers if they have the same block size and address. If there is a requirement for volume sharing, for example in a clustering environment, then a volume can be a member of multiple volume groups, or multiple servers can be in the same volume group. In both cases, data integrity is the responsibility of host software.

To create open systems volume groups, perform the following steps:

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Open systems** → **Volume groups**. On the Volume groups main page, in the **Select storage unit** drop-down list, select a storage unit.
2. From the **Select Action** list, select **Create**, and then click **Go**. The Create New Volume Group – Define volume group properties page is displayed.
3. Define the properties. In the **Volume Group Nickname** field, you can either keep or modify the default nickname. In **Accessed by host types**, select the appropriate host type. All other valid compatible host types will be automatically selected. Both the **Nickname** and **Accessed by host types** fields are required. After you define the properties, click **Next**. The Create Volume Group – Select host attachments page is displayed.
4. In Create Volume Group – Select host attachments page, you can optionally select the host attachment for the specified volumes. After selecting the host attachment, click **Next**. The Create Volume Group – Select volumes for groups page is displayed.

Note: If there are no host attachments available for selection, you can create one by clicking the **Create new host attachment** button. When you have finished creating a host attachment, you will be returned to this page, and the host attachment that you created will be available for selection.

5. In the Create Volume Group – Map host attachment page, select the I/O ports that the host attachment uses to access the volume group. Click **Next** to continue. The Create Volume Group – Select volumes for groups page is displayed.
6. Select the volumes for the volume group. At least one volume must be selected in the table. The volumes listed should be compatible with the host types selected in the previous two pages. After selecting the volumes for the volume group, click **Next**. The Create Volume Group – Verification page is displayed.

Note: If there are no volumes available for selection, you can create one by clicking the **Create new open systems volume** button. When you have finished creating a volume, you will be returned to this page, and the volume that you created will be available for selection.

Tip: If you created volumes with the same nickname prefix, you can change the view to show only those volumes that you want to select either by sorting the nickname column or by creating a filter.

7. Use the Create Volume Group – Verification page to review the attributes and verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and to specify the correct values. Otherwise, click **Finish** to complete the volume group creation process.

Applying a configuration (simulated only)

Complete this task to select, authenticate, and apply a storage unit configuration.

Ensure that you have obtained and applied an Operating Environment License with a maximum storage amount that exceeds the amount of storage that you have configured in your simulated configuration.

The source storage unit must have defined a logical configuration and hosts before you can apply a configuration.

To apply a configuration in simulated mode, perform the following steps:

1. In the navigation, select **Simulated manager** → **Manage hardware** → **Storage units**.
2. From the **Select Action** list, select **Apply configuration...**, and then click **Go**. The Apply Configuration — Select application method page is displayed.
3. In the Select application method page, specify the method with which to apply the configuration.
 - a. If you choose **Select from a list of storage complex**, the Select storage complex page of this wizard is displayed with the storage complexes in your simulated environment. Click **Next** to continue and go to step 3.
 - b. If you choose **Import new storage complex** the Import storage complex wizard is displayed. Once you finish with the wizard, the Select storage unit page is displayed with the storage units from the imported storage complex. You must be connected to the network to import the storage complex. Click **Next** to continue and go to step 5.
 - c. If you choose the **Apply configuration without importing storage complex** option, the Authenticate page is displayed. Click **Next** to continue and go to step 4.
4. Use the Select storage complex page to connect directly to a storage unit. Click **Next** to continue. The Authenticate page is displayed.
5. In the Authenticate page, connect to and authenticate a storage complex by defining management console properties. You must provide a user ID and password to complete the authentication. Click **Next** to continue. The Select storage unit page is displayed.
6. On the Select storage unit page, you connect directly to a storage unit. Specify the values as appropriate and click **Next**.
7. Use the Verification page to review the attributes and verify that they are correct.
8. If the attributes and values are not correct, click **Back**, as appropriate, to return and specify the correct values. Otherwise, click **Finish** to complete the apply configuration process.

If you are creating a new logical storage configuration, return to Chapter 3, “Creating a custom logical storage configuration,” on page 9 and complete the rest of the steps.

Enabling host system I/O

Complete this task to perform the high-level steps that are required to enable your host to send I/O to the DS6000.

You must complete your logical storage configuration before you attempt to send host I/O to the DS6000.

This task contains the high-level tasks that you must complete before you can successfully send host I/O to your DS6000. Instructions and details for each individual step can be found in the host system documentation or the *IBM System Storage DS6000 Host Systems Attachment Guide*.

1. Verify and update the supported fabric switch firmware.
2. Create your storage area network (SAN) zones and assign them to the appropriate zone sets.
3. Verify and update the host adapter firmware on the host system.
4. Install, verify, and update the host failover driver (SDD) on the host system.
5. If necessary, reconfigure or reboot the host system.
6. Verify that the host system can access the DS6000 logical volumes that are assigned to it.

If you are creating a new logical storage configuration, return to Chapter 3, “Creating a custom logical storage configuration,” on page 9 and complete the rest of the steps.

Chapter 4. Managing logical storage configurations

The topics in this category present information related to working with the various elements of an existing logical configuration, such as arrays, ranks, extent pools, and volumes. Information on the pages that are associated with these tasks can be found in the Reference section.

Viewing status (real-time only)

Complete this task to view the status of the elements in the active page.

1. In the **Select Action** drop-down list, select **Status....**
2. Next, select **Go**. The status page is displayed.
3. Click on the **Refresh** button to have the browser refresh and display current data.
4. Terminate the status page by clicking on the **OK** button or on the x in the page bar.

Viewing and modifying properties

Complete this task to access properties pages so that you can review or modify the properties.

1. In the main page of the target subject area (for example, Host Systems), select an item in the table.
2. In the **Select Action** drop-down list, select **Properties** and then **Go**. The properties page for the selected item is displayed.
3. If the properties page is solely informational, the only button option is **OK**. Click it to close the page. If the properties page has attributes that you can modify, the button options are **OK** or **Apply**. Click **OK** to apply the changes and close the page. Click **Apply** to apply the changes and leave the page open.

Viewing array disk drive module information

Complete this task to view array disk drive module (DDM) information.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Arrays**. On the Arrays main page, in the **Select storage unit** drop-down list, select the storage unit that contains the array that you want to view.
2. Select an array. In the **Select Action** drop-down list, select **Array DDMs**, and then click **Go**. The Array DDMs page is displayed with the DDM information for the selected array.
3. When you are finished reviewing the DDM information, click the **OK** button to close the page.

Adding arrays to ranks

Complete this task to add arrays to ranks and to specify the storage type.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Arrays**. On the Arrays main page, in the **Select storage unit** drop-down list, select the storage unit that contains the array that you want to add.

2. Select the array that you want to add to a rank. In the **Select Action** drop-down list, select **Add Array To Rank**, and then click **Go**. The Add Array to Rank page is displayed.
3. The Rank number is supplied by default. Accept or change the storage type.
4. Click the **OK** button to apply the array.

Modifying ranks

Complete this task to modify ranks and their specifications with regard to extent pools and storage types.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Ranks**. On the Ranks main page, in the **Select storage unit** drop-down list, select the storage unit that contains the rank that you want to modify.
2. Select a rank. In the **Select Action** drop-down list, select **Modify...**, and then click **Go**. The Modify Rank – Select extent pool page is displayed.
3. Select the extent pool that you want to include for this rank. If you need to create a new extent pool, click the **Create new extent pool** button to add a new extent pool. When you have finished creating an extent pool, you will be returned to this page, and the extent pool that you created will be available for selection. Click **Next**. The Modify Rank — Verification page is displayed.
4. On the Modify Rank — Verification page, review the attributes and values to verify that they are correct.
5. If the attributes and values are not correct, click < **Back** to return and specify the correct values. Otherwise, click **Finish** to complete the rank modification process.

Repairing a rank (real-time only)

Complete this task to repair a failed or inaccessible rank.

1. In the navigation, select **Real-time manager** → **Configure storage** → **Ranks**.
2. In the **Select Action** drop-down list, select **Status...**, and then click **Go**. The status page is displayed. If the selected rank has a status of Failed or Inaccessible, the **Repair** button is displayed. You can view a description of the event and the risk involved with the repair.

Note: If you repair the rank, all data will be lost.

3. To repair the rank, click **Repair**. All data is lost.

Adding ranks to extent pools

Complete this task to select or create the extent pools that you want to associate with ranks.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Ranks**. On the Ranks main page, in the **Select storage unit** drop-down list, select the storage unit that includes the rank that you want to assign to an extent pool.
2. Select one or more of the unassigned ranks of the same storage type. In the **Select Action** drop-down list, select **Add to Extent Pool**, and then click **Go**. The Ranks – Add to extent pool page is displayed.

3. Select your target extent pool. If you need to create a new extent pool, click the **Create new extent pool** button. When you have finished creating an extent pool, you will be returned to this page, and the extent pool that you created will be available for selection.
4. Click **OK** to complete the process and close the page.

Removing ranks from extent pools

Complete this task to remove a rank from an extent pool if the rank has no extents being used.

This option is available in the **Select Action** drop-down list only when one or more ranks that have the following characteristics are selected in the table:

1. The rank is currently assigned to a extent pool.
2. The rank does not have any extents used.
1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Ranks**. On the Ranks main page, in the **Select storage unit** drop-down list, select the storage unit that includes the rank that you want to remove from an extent pool.
2. Select one or more of the assigned ranks that you want to remove. In the **Select Action** drop-down list, select **Remove from extent pool** and then click **Go**. A confirmation dialog box is displayed.
3. Select **OK** to complete the removal process.

Modifying extent pools

Complete this task to modify extent pools.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Extent pools**. On the Extent pools main page, in the **Select storage unit** drop-down list, select the storage unit that includes the extent pool that you want to modify.
2. Select an extent pool. In the **Select Action** drop-down list, select **Modify...**, and then click **Go**. The Modify Extent Pool – Define properties page is displayed.
3. Modify the nickname. Click **Next**. The Modify Extent Pool — Select ranks page is displayed.
4. Select a rank, or create a new rank. After you have selected or created a rank or ranks, click **Next**. The Modify Extent Pool — Reserve storage page is displayed.
5. Specify the percent of reserve storage to modify or leave it unchanged. Click **Next**. The Modify Extent Pool — Verification page is displayed.
6. In Verification, review the attributes and values to verify that they are correct.
7. If the attributes and values are not correct, click **Back** as appropriate to return and then specify the correct values. Otherwise, click **Finish** to complete the extent pool modification process.

Modifying open systems volume groups

Complete this task to modify open systems volume group properties and parameters.

The Modify option is available in the **Select Action** drop-down list only after you have selected one of the groups in the table.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Open systems** → **Volume groups**. On the Volume groups main page, in the **Select storage unit** drop-down list, select the storage unit that contains the volume group that you want to modify.
2. Select the volume group that you want to modify. In the **Select Action** drop-down list, select **Modify...** The Modify Volume Group — Define volume group properties page is displayed.
3. Accept or modify the properties according to the following conditions: If you modify the host types that can access this volume group, then the host attachments in the Selected host attachments page might also require modification. You must make changes that are compatible with the current definition of the volume group. For example, you cannot add host types with differing block size and addressing methods, and you cannot remove all the host types from the selection. Valid modifications include adding other host types that are compatible with the current volume group definition and removing host types (as long as at least one remains).
4. Click **Next** to continue. The Modify Volume Group — Select host attachments page is displayed.
5. Accept or modify values. You can specify a different nickname from this page. If you change the host attachment, then the Selected volumes might also require modification. If you changed from a particular host type to another, namely from AS/400 to any other open systems host (or vice versa), then the subsequent pages require modification. Finish will be disabled until the user makes all the required changes.
6. Click **Next** to continue. The Modify Volume Group – Select volumes for groups page is displayed.
7. In the Modify Volume Group – Map host attachment page, accept or modify the I/O ports that the host attachment uses to access the volume group. Click **Next** to continue. The Modify Volume Group – Select volumes for groups page is displayed.
8. Accept or modify the volumes for the volume group. Select the volumes that you want to use in this volume group. You can use the **Create new volumes** button to create a new volume. After you create the volume and the table resets, the new volume is available for selection.

Tip: If you created volumes with the same nickname prefix, you can change the view to show only those volumes that you want to select either by sorting the nickname column or by creating a filter.

9. Click **Next** to continue. The Modify Volume Group — Verification page is displayed.
10. Use the Verification page to review the attributes and verify that they are correct.
11. If the attributes and values are not correct, click **Back** as appropriate to return and then specify the correct values. Otherwise, click **Finish** to complete the volume group modification process.

Modifying zSeries volumes

Complete this task to modify the properties and attributes of established zSeries volumes.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **zSeries** → **Volumes – zSeries**. On the Volumes – zSeries main page, in the **Select storage unit** drop-down list, select the storage unit

that includes the volume that you want to modify. In the **Select LCU** drop-down list, select the zSeries LCU that includes the volume or volumes that you want to delete.

2. In the **Select Action** drop-down list, select **Modify...**, and then click **Go**. The Modify zSeries volume — Select extent pool page is displayed. You cannot modify values here.
3. Click **Next**. The Define base volume characteristics page is displayed. You cannot modify values here.
4. Click **Next**. The Define base volume properties page is displayed. You cannot modify values here. Refer to the page help for the details on the possible pages to be displayed and the various fields that they contain.
5. Click **Next**. The Properties of the base volume page is displayed. You cannot modify values here.
6. Click **Next**. The Create volume nicknames page is displayed. Check the Generate a sequence of nicknames based on the following box to specify a sequence of volume names for the quantity created. Specify a prefix or suffix as appropriate.
7. Click **Next** to continue. The Define base alias page is displayed. You can only remove its assignment from the base volumes.
8. Click **Next** to continue. The Verification page is displayed.
9. Use the Verification page to review the attributes and verify that they are correct.
10. If the attributes and values are not correct, click **Back** as appropriate to return and specify the correct values. Otherwise, click **Finish** to complete the volume creation process.

Modifying LCUs

Complete this task to modify the attributes and properties of established logical control units (LCUs).

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **zSeries** → **LCUs**. On the LCUs main page, in the **Select storage unit** drop-down list, select the storage unit that includes the LCU that you want to modify.
2. In the **Select Action** drop-down list, select **Modify...**, and then click **Go**. The Modify LCU properties page is displayed.
3. Define the parameters to modify the selected LCU.
 - a. Select the LCU ID of the LCU that you want to modify.
 - b. You can change the default SSID by highlighting and typing over it.
 - c. Specify the LCU types and establish the timeout times in seconds.
 - d. Select a mode and, as necessary, establish the consistency group timeout in seconds.
4. Once you have modified the values, click **OK** to complete the process and close the page.

Chapter 5. Deleting a logical storage configuration

Follow these steps to delete or remove a logical storage configuration for open systems storage units. This applies to entire configurations including, ranks, arrays, volumes, and extent pools.

Recommendation: Because this procedure deletes the host and logical storage configuration from the applicable storage unit you can, prior to performing this task, export and save the current configuration so that it can be restored.

When you are deleting a logical storage configuration, individually delete the hosts and volume group.

To delete a logical configuration, perform the following steps:

1. Delete host access to the volume groups that you want to remove. To remove the host port IDs use the subtask in Deleting host systems.
2. Delete the volume group. To delete volume groups, you must complete the subtask in Deleting open system volume groups.
3. Delete the array. To delete an array, you must complete the subtask in Deleting arrays.

Note: Any ranks that make up an extent pool, along with the volumes within the extent pool, are automatically deleted if the rank is assigned to the array when you delete the array.

4. Delete any unassigned ranks using the subtask in Deleting ranks.
5. Delete the extent pools, if they were not deleted by prior delete actions. To delete the extent pools for open systems volumes, you must complete the subtask in Deleting extent pools.

Verify that no logical configuration exists for the applicable storage unit.

Deleting host systems

Complete this task to remove host systems from the configuration.

1. In the navigation, select **Real-time manager or Simulated manager** → **Manage hardware** → **Host systems**.
2. Select the storage complex from the drop-down list.
3. In the table on the Host systems - Main page, select the host system that you want to delete.
4. In the **Select Action** drop-down list, select **Delete** and then **Go**. A confirmation dialog box is displayed.
5. Confirm the action to complete the deletion of the selected host systems by clicking **Continue**.

Deleting open systems volume groups

Complete this task to delete open systems volume groups, releasing any assigned volumes for assignment to other volume groups.

If the number of volumes in this volume group is not zero, the volumes that are assigned to this volume group go into an unassigned mode where they are eligible for assignment to other volume groups.

Note: The volume group is the link between the open systems volumes and the host attachments. If you delete a volume group that is attached to a host attachment identifier, the host attachment is not able to access the volumes.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Open systems** → **Volume groups**. On the Volume groups main page, in the **Select storage unit** drop-down list, select the storage unit that contains the volume group that you want to delete.
2. Select the volume group or groups that you want to delete.
3. In the **Select Action** drop-down list, select **Delete**, and then click **Go**. A confirmation dialog box is displayed.
4. Confirm the action to complete the deletion of the selected volume groups.

Deleting arrays

Complete this task to delete arrays with and without volumes using storage.

If the arrays you are deleting are assigned to volume groups, you must first remove the arrays from the volume groups before you can delete them.

If the selected arrays have volumes using storage on the array, an **Attention** message indicates the volumes that have storage configured on the arrays. The message also indicates that the volumes will be deleted as part of the array deletion process.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Arrays**. On the Arrays main page, in the **Select storage unit** drop-down list, select the storage unit that contains the array that you want to delete.
2. Select the array or arrays that you want to delete. In the **Select Action** drop-down list, select **Delete**, and then click **Go**.

Note: If the array that you delete has volumes using storage on it, these volumes will be deleted along with the array. You will see a message displaying the volume numbers and nicknames that will be deleted and be asked whether you want to continue with or cancel the deletion process.

A confirmation message is displayed. If the volumes are assigned, regardless of any storage used, the confirmation indicates that the array will be deleted.

3. Click **OK** to complete the deletion of the selected arrays.

Deleting ranks

Complete this task to delete ranks with or without volumes using storage.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Ranks**. On the Ranks main page, in the **Select storage unit** drop-down list, select the storage unit that contains the rank that you want to delete.
2. Select the rank or ranks that you want to delete. In the **Select Action** drop-down list, select **Delete**, and then click **Go**. A confirmation dialog box is displayed.

Note: If the ranks you selected have volumes using storage on the array, an **Attention** message lists the volumes that have storage configured on the ranks. The message also indicates that the volumes will be deleted as part of the rank deletion process.

3. If the **Attention** message is displayed, click on **Continue** or **Cancel**. If you click on **Cancel**, you are returned the Ranks main page. If you click on **Continue**, go to the next step.
4. After you click on **Continue**, a confirmation dialog is displayed. If the selected ranks are unassigned or do not have any extents allocated on the ranks, confirming the dialog box deletes the selected ranks. The array assigned to this rank goes into an unavailable status. In an unassigned status, the array becomes eligible for assignment to another rank.
5. Confirm the dialog box to complete the deletion of the selected ranks and arrays.

Deleting open systems volumes

Complete this task to delete open systems volumes from the storage unit.

You must make a selection in the table to enable this option.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Open systems** → **Volumes – Open systems**. On the Volumes – Open systems main page, in the **Select storage unit** drop-down list, select the storage unit that contains the volume that you want to delete.
2. Select the volume or volumes that you want to delete.
3. In the **Select Action** drop-down list, select **Delete** and then **Go**. A confirmation dialog box is displayed.
4. Confirm the action to complete the deletion of the selected volumes.

Deleting extent pools

Complete this task to delete extent pools with or without volumes using storage.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **Extent pools**. On the Extent pools main page, in the **Select storage unit** drop-down list, select the storage unit that includes the extent pool that you want to delete.
2. On the Extent pools main page, select the extent pools that you want to delete.
3. In the **Select Action** drop-down list, select **Delete**, and then click **Go**. A confirmation dialog box is displayed.
4. If the selected extent pools have volumes configured in them, an **Attention** message indicates the count of volumes that have storage configured in the extent pools. The message also indicates that these volumes are deleted as part of the extent pool deletion process.
5. If the **Attention** message is displayed, click on **Continue** or **Cancel**. Clicking on **Cancel** returns you to the Ranks page. If you click on **Continue**, go to the next step.
6. After you click on **Continue**, a confirmation dialog is displayed. If you confirm the dialog, any ranks that are assigned to this extent pool are released into an unassigned configuration state. These ranks become eligible for assignment to another extent pool.
7. Confirm the dialog box to complete the deletion of the selected extent pools.

Deleting zSeries volumes

Complete this task to delete zSeries volumes and any of their associated aliases that are not aliases for other volumes.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **zSeries** → **Volumes – zSeries**. On the Volumes – zSeries main page, in the **Select storage unit** drop-down list, select the storage unit that includes the volume or volumes that you want to delete. In the **Select LCU** drop-down list, select the zSeries LCU that includes the volume or volumes that you want to delete.
2. Select the zSeries volumes that you want to delete. This process deletes the zSeries volumes and any associated aliases that are not aliases for other volumes.
3. In the **Select Action** drop-down list, select **Delete**, and then click **Go**. A confirmation dialog box is displayed.
4. Confirm the action to complete the process.

Deleting logical control units (LCUs)

Complete this task to delete LCUs and any of their associated volumes.

If the selected LCUs have volumes defined in them, an **Attention** message indicates that there are volumes configured in the LCUs. The message also indicates that these volumes are deleted as part of the LCU deletion process.

1. In the navigation, select **Real-time manager or Simulated manager** → **Configure storage** → **zSeries** → **LCUs**. On the LCUs main page, in the **Select storage unit** drop-down list, select the storage unit that includes the LCU or LCUs that you want to delete.
2. Select the LCU or LCUs that you want to delete. In the **Select Action** drop-down list, select **Delete**, and then click **Go**. A confirmation dialog box is displayed.
3. If the **Attention** message is displayed, click on **Continue** or **Cancel**. Clicking on **Cancel** returns you to the LCUs page. If you click on **Continue**, go to the next step.
4. After you click on **Continue**, a confirmation dialog is displayed and, if confirmed, the selected LCUs and volumes are deleted. If the selected LCUs do not have volumes configured in them, a confirmation dialog is invoked and, if confirmed, the characteristics of the selected LCUs are deleted.
5. Confirm the dialog box to complete the deletion of the selected LCUs and volumes.

Chapter 6. DS6000 host attachment

The DS6000 storage unit provides a variety of host attachments so that you can consolidate storage capacity and workloads for open-systems hosts, S/390® hosts, and eServer™ zSeries® hosts.

You can configure the storage unit for any of the following system adapter types and protocols:

- Fibre-channel adapters, for support of fibre-channel protocol (FCP) and fibre connection (FICON®) protocol

General requirements for attaching a host system

Before you can attach a DS6000 to a host system, you must review a list of the general requirements that pertain to all hosts. You must also review the specific host requirements described in the introduction to each specific host.

Perform the following steps before you attach any host system to a DS6000:

1. Check the DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most up-to-date details about the release level for your operating system and adapter cards.
2. Retrieve a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>.
3. Ensure that you can reference the following documentation:
 - The *IBM System Storage DS6000 Installation, Troubleshooting, and Recovery Guide* from the System Storage™ CD that you receive with the storage unit.
 - The IBM System Storage DS6000 Information Center at <http://publib.boulder.ibm.com/infocenter/ds6000ic/index.jsp>.
4. Install the storage unit by using the procedures in the *IBM System Storage DS6000 Installation, Troubleshooting, and Recovery Guide*.
5. Check the LUN limitations for your host system and verify that there are enough adapters that are installed in the server to manage the total LUNs that you want to attach. See LUN considerations for fibre-channel attachment.
6. Review the "I/O adapter features" section of the *IBM System Storage DS6000 Introduction and Planning Guide* for detailed host adapter configuration rules and for information on host adapter and cable feature codes, host adapter balancing for optimal performance, and supported cabling distances.
7. Use the DS Storage Manager or the DS CLI to define the host and I/O port configurations. Make sure that you define the worldwide port names for fibre-channel ports. See DS8000 Interfaces for more information on the DS Storage Manager and DS CLI interfaces.
8. Install the adapter driver that came with your host adapter or use the steps defined in "Downloading and installing a host adapter driver" on page 30 to download and install an updated adapter driver.

Downloading and installing a host adapter driver

This section describes the steps you can follow to download and install an AMCC/JNI, Emulex, HP, Netfinity, QLogic or Sun host adapter driver.

Perform the following steps to download and install a host adapter driver:

Note: You must download relevant vendor documentation for the driver that you select to correctly install and configure the host adapter.

1. Verify in the DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> that the host adapter you want to use is compatible with the DS6000, your host, and your host operating system.
2. Retrieve a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss..>
3. Download the adapter driver from the adapter vendor Web site. Table 1 provides vendor URLs and specific download steps. If the steps to locate the download page are not current, use the search function on the vendor Web site to locate the appropriate download for your host adapter.

Table 1. Host adapter driver download steps by vendor

Host Adapter Vendor	URL	Steps to locate download page
AMCC/JNI	http://www.amcc.com	<ol style="list-style-type: none">1. Select Drivers/Downloads from the top menu bar.2. Select FC HBA/OEM Files.3. Select the IBM button.4. Scroll down or use the search function of your browser to locate your adapter.5. Download the appropriate files for your adapter.
Emulex	http://www.emulex.com/ts/docoem/ibm/index.htm	<ol style="list-style-type: none">1. Scroll down or use the search function of your browser to locate your operating system in the list of drivers.2. Select the driver from the list.3. Download the driver and the Configuration Utility tool that is provided with the driver.

Table 1. Host adapter driver download steps by vendor (continued)

Host Adapter Vendor	URL	Steps to locate download page
Hewlett-Packard	http://www.hp.com	<ol style="list-style-type: none"> 1. Enter the name of the host adapter in the Search field at the bottom of the Hewlett-Packard home page. 2. In the search results, click the link for the driver for your operating system. 3. Click the Download button to download the adapter driver. 4. Return to the search results from step 1 and review the links to documentation, such as installation requirements and release notes.
IBM (Netfinity)	http://www-1.ibm.com/servers/storage/support/disk	<ol style="list-style-type: none"> 1. Select your storage unit from the Enterprise storage servers list. 2. Click the HBA interoperability search tool under Technical Resources in the Troubleshooting tab. 3. Select the appropriate options for your product and operating system and click the Submit button at the bottom of the page. 4. Locate the section for the current version of the driver and firmware that you want, and click View Details. 5. Click the Driver Level that you want. 6. Click Continue to leave the IBM System Storage support Web site. 7. Click one of the Download options for the feature code that you want.

Table 1. Host adapter driver download steps by vendor (continued)

Host Adapter Vendor	URL	Steps to locate download page
QLogic	http://support.qlogic.com/support/oem_ibm.asp	<ol style="list-style-type: none"> 1. Click the DS6000 link under the IBM System Storage Supported Software list. 2. Scroll down or use the search function of your browser to locate the table for your adapter and click the Download link to download the version you want for the driver for your operating system.
Sun	http://www.sun.com/storage/san	<ol style="list-style-type: none"> 1. Scroll down to the section titled Get the Software. 2. Locate the current driver in the list and click the appropriate link. 3. Type the Username and Password in the fields and click begin. If you do not have a username and password, you may begin the registration process and return to this page when you have acquired a username and password. Note: Sun only grants usernames and passwords to customers who have purchased maintenance contracts. 4. Select Accept to agree to the license agreement (required). 5. Select the appropriate links and click on them to begin the download process.

4. Follow the installation instructions from the vendor to install the host adapter driver. Some installation instructions might be included in the readme files that are included in the compressed download, but some installation documentation might have to be downloaded separately. Return to the vendor URL that is listed in Table 1 on page 30 to locate installation and configuration instructions.
5. Locate the chapter for your host in the *IBM System Storage DS6000 Host Systems Attachment Guide*, and follow any additional driver configuration instructions.

Host attachment general considerations

This section contains general conceptual overviews on key topics for host attachment.

IBM System Storage Multipath Subsystem Device Driver

The IBM System Storage Multipath Subsystem Device Driver (SDD) supports redundant connections between a disk storage server and a host server.

The SDD resides in the host server with the native disk device driver and takes advantage of redundant storage configurations to provide single-point failure protection, high data availability, and storage concurrent firmware download. It also provides a dynamic I/O load-balancing algorithm to reach the best I/O performance.

The SDD provides the following functions:

- Enhanced data availability
- Automatic path failover and failback
- Dynamic I/O load-balancing across multiple paths
- Multiple path selection policies for the host system
- Concurrent download of licensed machine code

The term *SDD* represents both the SDD as the established subsystem multipath device driver and as the multipath I/O (MPIO) path control module, depending upon the operating system. Table 2 provides examples.

Table 2. Examples of SDD MPIO path control modules

Operating system	Description	Acronym
AIX®	SDD Path Control Module	SDDPCM
Linux®	Device Mapper	DM-Multipath
Windows®	SDD Device Specific Module	SDDDSM

As new operating systems are added, the SDD will be named according to the operating system.

For more information about the SDD, see <http://www.ibm.com/systems/support/storage/software/sdd>.

Attachment restrictions for Copy Services

This section describes attachment restrictions for FlashCopy® SE and Remote Mirror and Copy.

When you copy a source volume to a target volume with FlashCopy SE or Remote Mirror and Copy, the source and target volumes must be on different host systems to enable concurrent read/write access of both volumes.

A copy operation with the source and target volume on the same host system creates a target volume with the same identification as the source volume. The result is that the host system sees two identical volumes. When the copy operation

creates the same identification for the target volume as for the source volume, you are not able to distinguish one from the other. Therefore, you might not be able to access the original data.

Note: You cannot create a host target on a single Novell NetWare host system. For Novell NetWare, the target volume must be attached to a second Novell NetWare host system.

The target volume and the source volume can be on the same host system for a FlashCopy SE or Remote Mirror and Copy operation only under the following conditions:

- For AIX, when the host system is using a logical volume manager (LVM) with **recreatevg** command support.
- For AIX and Sun, when the host system is *not* using a logical volume manager.
- For HP using fibre-channel, when the host is using LVM with the **vfchigid -f** command.
- For any host system, when the host system can distinguish between a source and a target volume that have the same identification.

DS6000 fabric zoning

Creating zones for your host attachment port connections enables you to manage traffic and establish redundancy that eliminates single points of hardware failure.

Depending on your environment, you can benefit from isolating traffic as much as possible in your storage area network. Networks with a large amount of volumes and heavy host traffic can benefit the most. Implementing single-initiator, single-target zoning allows you to isolate traffic for each port. Single-initiator, single-target zoning creates small zones within the fabric with only two zone members (ports). The zone consists of one target (a storage unit port), and one initiator (a host system port). The key benefit of single-initiator, single-target zoning is traffic isolation or masking.

Figure 1 on page 35 displays two host systems that are attached to a DS6800 through two fabric switches. Through the switches, each of the host systems is connected to both of the DS6800 processor cards (PC-0 and PC-1) using two host adapters in each host system (HA-1 and HA-2). Traffic isolation and masking occur by creating a zone for each host system port.

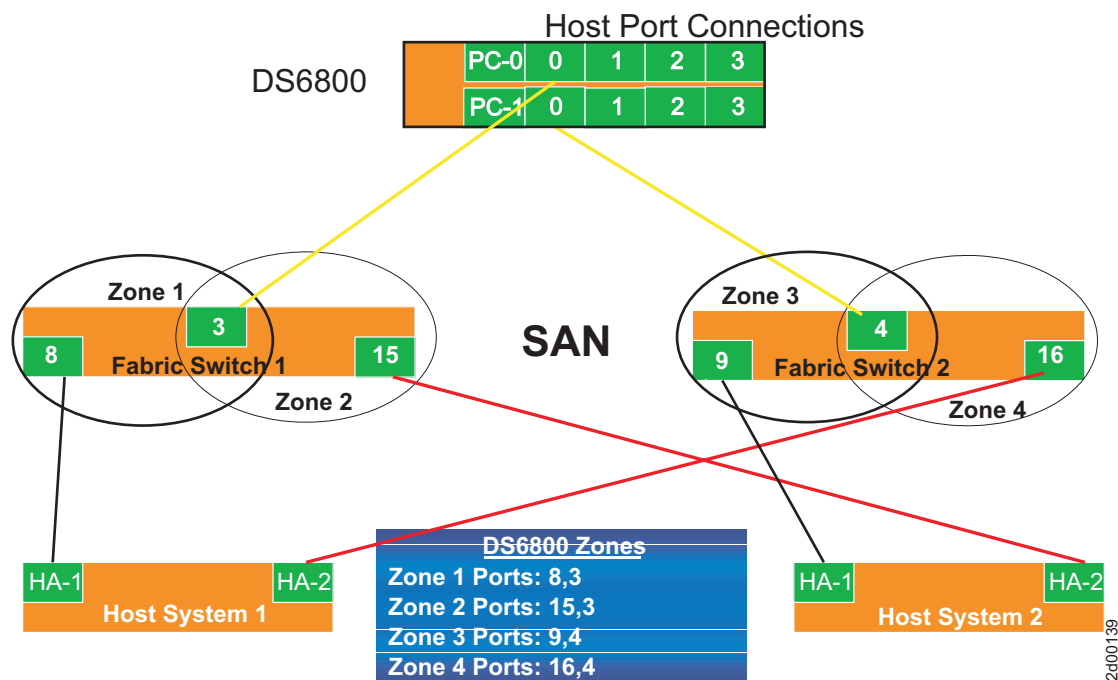


Figure 1. DS6000 zoning example

Apple Macintosh host attachment

This chapter lists the general requirements, supported hosts, operating system and adapters for you to attach an Apple Macintosh host to a DS6000 using fibre-channel adapters.

Before you attach a host, review the following information:

- The "General requirements for attaching a host system" on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

The host type selection is Intel Linux for the Apple server attachment.

Fujitsu PRIMEPOWER host attachment

This chapter describes how you attach a Fujitsu PRIMEPOWER host to a DS6000 using a fibre-channel Emulex adapter.

This chapter contains the following host attachment topics:

- Locating the WWPN for a Fujitsu PRIMEPOWER host
- Installing the Emulex adapter driver for a PRIMEPOWER host system
- Parameter settings for the Emulex LP9002L adapter
- Setting parameters for Emulex adapters
- Configuring host device drivers for PRIMEPOWER

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, switches and microcode
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Note: If you are installing the fibre-channel adapter for the first time, you must specify the correct topology. You must also select the appropriate device mapping driver.

Locating the WWPN for a Fujitsu PRIMEPOWER host

Complete this task to locate the WWPN for a PRIMEPOWER host system adapter.

1. After you install the adapter and you restart the host system, view the `/var/adm/messages` file.
2. Search for the line that contains the `lpfcx: Fibre Channel WNN` phrase, where `x` is the adapter number 0, 1, and so on.

Installing the Emulex adapter for a PRIMEPOWER host

This section describes how you configure an Emulex adapter card and driver for a PRIMEPOWER host.

Single- and dual-port, fibre-channel interfaces with an Emulex adapter card support the following public and private loop modes:

- Target
 - Public initiator
 - Private initiator
 - Target and public initiator
 - Target and private initiator
1. Record the IEEE number that is printed on the Emulex adapter card. You can use the IEEE number to determine the WWPN.
 2. Install the Emulex adapter card in your PRIMEPOWER host.
 3. Follow the steps to download and install an adapter driver in “Downloading and installing a host adapter driver” on page 30.

Parameter settings for the Emulex LP9002L adapter

You can use these recommended configuration settings for your Emulex adapter on a PRIMEPOWER host system.

Table 3 on page 37 provides a list of recommended parameters and setting for the Emulex LP9002L adapter.

Table 3. Recommended configuration file parameters for the Emulex LP9002L adapter

Parameters	Recommended settings
automap	1: Default. SCSI IDs for all FCP nodes without persistent bindings are automatically generated. If new FCP devices are added to the network when the system is down, there is no guarantee that these SCSI IDs will remain the same when the system is restarted. If one of the FCP binding methods is specified, then automap devices use the same mapping method to preserve SCSI IDs between link down and link up. If no bindings are specified, a value of 1 forces WWNN binding, a value of 2 forces WWPN binding, and a value of 3 forces DID binding. If automap is 0, only devices with persistent bindings are recognized by the system.
fcg-on	1: Default. Turn on FCP.
lun-queue-depth	30: The default value that the driver uses to limit the number of outstanding commands per FCP LUN. This value is global, affecting each LUN recognized by the driver, but can be overridden on a per-LUN basis. You might have to configure RAID using the per-LUN tunable throttles.
no-device-delay	0: Default. Implies no delay whatsoever. 1: Recommended. 2: Setting a long delay value might permit I/O operations to build up, each with a pending timeout, which could result in the exhaustion of critical PRIMEPOWER kernel resources. In this case, you might see a fatal message such as, PANIC: Timeout table overflow.
network-on	0: Default. Recommended for fabric. Do not turn on IP networking. 1: Turn on IP networking.
scan-down	0: Recommended. Causes the driver to use an inverted ALPA map, effectively scanning ALPAs from high to low as specified in the FC-AL annex. 2: Arbitrated loop topology.
tgt-queue-depth	0: Recommended. The default value that the driver uses to limit the number of outstanding commands per FCP target. This value is global, affecting each target recognized by the driver, but can be overridden on a per-target basis. You might have to configure RAID using the per-target tunable throttles.
topology	2: Recommended for fabric. Point-to-point topology only. 4: Recommended for nonfabric. Arbitrated-loop topology only.
xmt-que-size	256: Default. Size of the transmit queue for mbufs (128 - 10240).
zone-rscn	0: Default. 1: Recommended for fabric. Check the name server for RSCNs. Setting zone-rscn to 1 causes the driver to check with the name server to see if an N_Port ID received from an RSCN applies. If soft zoning is used with Brocade fabrics, this parameter should be set to 1.

Setting parameters for Emulex adapters

This section provides instructions for setting parameters for Emulex adapters for your PRIMEPOWER host system.

1. Type `cd /etc` to change to the `/etc` subdirectory.

2. Backup the system file in the subdirectory.
3. Edit the system file and set the following parameters for servers with configurations that only use Emulex or QLogic adapters.

sd_io_time

This parameter specifies the timeout value for disk operations. Add the following line to the `/etc/system` file to set the `sd_io_time` parameter for the storage unit LUNs: `set sd:sd_io_time=0x78`

sd_retry_count

This parameter specifies the retry count for disk operations. Add the following line to the `/etc/system` file to set the `sd_retry_count` parameter for the storage unit LUNs: `set sd:sd_retry_count=5`

maxphys

This parameter specifies the maximum number of bytes you can transfer for each transaction. The default value is 12 6976 (124 KB). If the I/O block size that you requested exceeds the default value, the request is broken into more than one request. You should tune the value to the application requirements. For maximum bandwidth, set the `maxphys` parameter by adding the following line to the `/etc/system` file: `set maxphys=1048576 (1 MB)`

Note: Do not set the value for `maxphys` greater than 1048576 (1 MB). Doing so can cause the system to hang.

Configuring host device drivers for a PRIMEPOWER host

Complete this task to update the PRIMEPOWER driver configuration file. This procedure gives you access to target and LUN pairs that are configured on the storage unit.

Attention: Do not change or remove entries for preexisting devices in `kernel/drv/sd.conf`. Doing so can cause your system to become inoperable.

1. Type `cd /kernel/drv` to change to the directory.
2. Back up the `sd.conf` file in this subdirectory.
3. Edit the `sd.conf` file to add support for the target and LUN pairs that are configured on the host system.

Note: Do not add duplicate target and LUN pairs.

Figure 2 on page 39 shows the lines that you must add to the file to access LUNs 0 - 49 on target 0 for fibre-channel.

```

name="sd" class="scsi"
    target=0 lun=0;
name="sd" class="scsi"
    target=0 lun=1;
name="sd" class="scsi"
    target=0 lun=2;
name="sd" class="scsi"
    target=0 lun=3;
name="sd" class="scsi"
    target=0 lun=4;
name="sd" class="scsi"
    target=0 lun=5;
name="sd" class="scsi"
    target=0 lun=6;
name="sd" class="scsi"
    target=0 lun=7;
name="sd" class="scsi"
    target=0 lun=8;
name="sd" class="scsi"
    target=0 lun=9;
name="sd" class="scsi"
    target=0 lun=10;
.
.
.
name="sd" class="scsi"
    target=0 lun=48;
name="sd" class="scsi"
    target=0 lun=49;

```

Figure 2. PRIMEPOWER host example of `sd.conf` file entries for fibre-channel

Figure 3 shows the start lpfc auto-generated configuration.

Note: You will delete anything that you put within this auto-generated section if you issue the **pkgrm** command to remove the lpfc driver package. You might want to add additional lines to probe for additional LUNs or targets. Delete any lines that represent lpfc targets or LUNs that are not used.

```

name="sd" parent="lpfc" target=0 lun=0;
name="sd" parent="lpfc" target=1 lun=0;
name="sd" parent="lpfc" target=2 lun=0;
name="sd" parent="lpfc" target=3 lun=0;
name="sd" parent="lpfc" target=4 lun=0;
name="sd" parent="lpfc" target=5 lun=0;
name="sd" parent="lpfc" target=6 lun=0;
name="sd" parent="lpfc" target=7 lun=0;
name="sd" parent="lpfc" target=8 lun=0;
name="sd" parent="lpfc" target=9 lun=0;
name="sd" parent="lpfc" target=10 lun=0;
name="sd" parent="lpfc" target=11 lun=0;
name="sd" parent="lpfc" target=12 lun=0;
name="sd" parent="lpfc" target=13 lun=0;
name="sd" parent="lpfc" target=14 lun=0;
name="sd" parent="lpfc" target=15 lun=0;
name="sd" parent="lpfc" target=16 lun=0;
name="sd" parent="lpfc" target=17 lun=0;
name="sd" parent="lpfc" target=17 lun=1;
name="sd" parent="lpfc" target=17 lun=2;
name="sd" parent="lpfc" target=17 lun=3;

```

Figure 3. PRIMEPOWER host example of a start lpfc auto-generated configuration

4. Type either:

- **reboot -- -r** from the Open Windows window to shut down and restart the PRIMEPOWER host system with the kernel reconfiguration option.

- **boot -r** from the OK prompt after you shut down.

The fibre-channel adapters that are supported for attaching the storage unit to a PRIMEPOWER host are capable of full-fabric support. Ensure that all fibre-channel driver configurations include worldwide port name, worldwide node name, port ID, or host adapter binding of target LUN pairs.

Binding of target LUN pairs implements the PRIMEPOWER fibre-channel host adapter configuration file that is installed by the adapter software package. Refer to the manufacturer's adapter documentation and utilities for detailed configuration instructions.

You can tune fibre-channel host adapter configuration files for host system reliability and performance.

Hewlett-Packard Server (HP-UX) host attachment

This chapter describes how you attach a Hewlett-Packard (HP-UX) host system to a DS6000 with fibre-channel adapters.

This chapter contains the following sections:

- Locating the WWPN for an HP-UX host
- Setting the queue depth for an HP-UX host
- Configuring clustering on an HP-UX host

Before you attach a host, review the following information:

- "General requirements for attaching a host system" on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> which provides current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>
- The *IBM System Storage Multipath Subsystem Device Driver User's Guide* at <http://www.ibm.com/systems/support/storage/software/sdd>

Notes:

1. To have failover protection on an open system, SDD allows a maximum of 32 paths, but requires a minimum of two paths.
2. Do not configure the same logical unit number (LUN) to be managed by both SDD and PVLINKS. Do not configure the same volume group with LUNs that are managed by both SDD and PVLINKS.
3. Due to limitations in the host operating system, LUNs greater than x'3FFF' are not supported. When you create or assign LUNs and volumes, only LUN and volume IDs less than x'3FFF' are supported. This limits the maximum number of volumes that are allowed for HP host types to 16384.

Locating the WWPN for an HP-UX host

Complete this task to locate the WWPN for a Hewlett-Packard Server host.

1. Go to the root directory of your HP-UX host.
2. Type `ioscan -fnC fc` | more for information on the fibre-channel adapters installed on the host.

The following is example output:

```
fc 0 0/2/0/0 td CLAIMED INTERFACE HP Tachyon XL2 Fibre Channel Mass Storage Adapter /dev/td0
fc 1 0/4/0/0 td CLAIMED INTERFACE HP Tachyon XL2 Fibre Channel Mass Storage Adapter /dev/td1
fc 2 0/6/2/0 td CLAIMED INTERFACE HP Tachyon XL2 Fibre Channel Mass Storage Adapter /dev/td2
```

3. Look under the description for the Fibre Channel Mass Storage adapter.
For example, look for the device path name `/dev/td1`.
4. Type: `fcmsutil /dev/td1 | grep world` where `/dev/td1` is the path.
The following is example output:

```
# fcmsutil /dev/td1 | grep World
N_Port Node World Wide Name = 0x50060b000024b139
N_Port Port World Wide Name = 0x50060b000024b138
(root@hpmain)/home/root# fcmsutil /dev/td0 | grep World
N_Port Node World Wide Name = 0x50060b000023a521
N_Port Port World Wide Name = 0x50060b000023a520
(root@hpmain)/home/root# fcmsutil /dev/td2 | grep World
N_Port Node World Wide Name = 0x50060b0000253a8f
N_Port Port World Wide Name = 0x50060b0000253a8e
(root@hpmain)/home/root#
```

Setting the queue depth for an HP-UX host

Complete this task to set the queue depth for a Hewlett-Packard Server host.

Before you set the queue depth, you must attach the host system to a DS6000. For instructions, see “Hewlett-Packard Server (HP-UX) host attachment” on page 40.

Procedure

1. Use the following formula to set the queue depth for all classes of HP-UX:
 $256 \div \text{maximum number of LUNs} = \text{queue depth}$

Note: Although this algorithm implies that the upper limit for the number of LUNs on an adapter is 256, HP-UX supports up to 1024 LUNs.

For optimum performance, monitor configurations with greater than 256 LUNs and adjust the queue depth.

2. Update the queue depth:
 - To update the queue depth by device level, type `scsictl -m queue_depth=21 /dev/rdisk/$dsksf`, where `/dev/rdisk/$dsksf` is the device node
 - To make a global change to the queue depth, use the HP System Administration Manager (SAM) and edit the kernel parameter so that it equals `scsi_max_qdepth`.

Configuring clustering on an HP-UX host

Complete this task to configure a storage unit for clustering on a Hewlett-Packard Server host that uses the MC/ServiceGuard clustering software.

The steps to configure MC/ServiceGuard with the storage unit are the same as the steps in the Hewlett-Packard high availability documentation at <http://www.docs.hp.com/en/B5140-90020/ch02s04.html>.

After you configure your host for normal operating system access, the storage unit acts as a normal disk device in the MC/ServiceGuard configuration. You can create volume groups that contain the volumes by using the Hewlett-Packard logical

volume manager. This method of disk management is more reliable, easier, and more flexible than whole-disk management techniques.

When you create volume groups, you can implement PV-Links, Hewlett-Packard's built-in multipathing software for high availability disks. You can also implement IBM Multipath Subsystem Device Driver (SDD), which is the multipathing software recommended by IBM.

Implementing PV-Links multipathing software

To use PV-Links multipathing software, perform the following steps:

1. Create the volume group, using the path to the volumes that you want as the primary path to the data. The following is an example showing the commands that you use to create the volume group:

```
mkdir /dev/vg15
mknod /dev/vg15/group c 64 0x030000
pvcreate /dev/rdisk/c1t0d0
pvcreate /dev/rdisk/c1t2d0
vgcreate /dev/vg15 /dev/dsk/c2t0d0
```

In this example, the new volume group is /dev/vg15 and the primary path to the data is /dev/dsk/c2t0d0.

2. Extend the volume group with the path to the volumes that are intended as alternate paths. The logical volume manager reads the label on the disk and knows that it is an alternate path to one of the volumes in the group. The logical volume manager labels the volume.

The following command creates an alternate path using the volume group that you created in Step 1 (/dev/vg15):

```
vgextend /dev/vg15 /dev/dsk/c3t0d0
```

In this example, /dev/dsk/c3t0d0 is the alternate path to the data.

3. To display the alternate path that you created in Step 2, type the following command:

```
vgdisplay -v /dev/vg15
```

For more information on these commands, see the online manpages `pvcreate(1M)`, `vgcreate(1M)`, and `vgdisplay(1M)`, or see the *HP-UX Reference (Manpages)* at http://www.docs.hp.com/en/hpuxman_pages.html.

Implementing SDD multipathing software

To use SDD multipathing software, perform the following steps:

1. Create the volume group using the vpath devices. The alternate paths of the volume are controlled by the vpath.
2. Extend the volume group of another vpath device only for multivolume groups.
3. Because vpath1 controls two device nodes c2t0d0 and c3t0d0, you only need to create a volume group with the vpath1 device.

With the SDD implementation, the utility `vpcluster` provides a convenient way for exporting and importing volume groups among the MC/Service Guards nodes.

Note: For more information about the SDD, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* at <http://www.ibm.com/systems/support/storage/software/sdd>.

HP AlphaServer or Integrity server OpenVMS host attachment

This chapter describes how you attach a Hewlett-Packard (HP) AlphaServer or Integrity server OpenVMS host system to a DS6000 with fibre-channel adapters.

This chapter contains the following sections:

- Confirming the installation of the OpenVMS operating system
- Installing and setting the mode for the KGPSA-xx adapter card in an AlphaServer OpenVMS host system
- Locating the WWPN for an OpenVMS host
- Managing devices for an Integrity server using the OpenVMS Extensible Firmware Interface
- Confirming fibre-channel switch connectivity for OpenVMS
- Defining a UDID for OpenVMS storage unit volumes
- Accessing the storage unit volumes from an OpenVMS host
- Guidelines for OpenVMS LUN 0
- Restrictions for OpenVMS fibre-channel connections

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>
- The HP document, *Guidelines for OpenVMS Cluster Configurations*, for a discussion about working with fibre-channel devices and recommendations for system parameters
- The HP document, *WWIDMGR User's Manual*, for a detailed description of the AlphaServer console commands that are used in this chapter. You can find this document on the Alpha Systems Firmware Update CD.
- The HP Technical Support Web site for Integrity server documentation.

Notes:

1. IBM supports cascaded switches in configurations up to a maximum of eight switches with a maximum of three interswitch hops for any path. Use two hops for normal operations and reserve the third hop for backup paths.
2. You do not need the Subsystem Device Driver because OpenVMS manages multipathing.
3. DS6000 host I/O ports no longer must be dedicated to an OpenVMS host. However, non-OpenVMS hosts logged into an OpenVMS-enabled port must abide by the OpenVMS mode behavior of no ABTS notifications.

Confirming the installation of the OpenVMS operating system

Complete this task to confirm the installation of the Hewlett-Packard OpenVMS operating system.

To confirm the operating system installation, perform the following steps:

1. To display the current version of the OpenVMS operating system that you have installed, type **show system**.

The following is example output:

```
$ show system
OpenVMS V7.3-2 on node DS20 5-FEB-2004 15:15:16.88 Uptime 1 01:45:09
```

2. Ensure that you have installed the most current version of patches and HP-recommended remedial kits. Some kits are dependent on other kits, so you must install some kits before you install another kit. See the documentation from Hewlett-Packard for more information.

3. To check the patches that are installed on your system, type **product show history**.

The following is example output:

```
$ product show history
```

PRODUCT	KIT TYPE	OPERATION	DATE AND TIME
DEC AXPVMS VMS731_LAN V6.0	Patch	Install	04-AUG-2004 15:34:30
DEC AXPVMS VMS731_UPDATE V1.0	Patch	Install	04-AUG-2004 15:28:37
DEC AXPVMS VMS731_PCSI V1.0	Patch	Install	04-AUG-2004 15:22:10
CPQ AXPVMS CDSA V1.0-2	Full LP	Install	04-AUG-2004 13:59:00
DEC AXPVMS DWMOTIF V1.2-6	Full LP	Install	04-AUG-2004 13:59:00
DEC AXPVMS OPENVMS V7.3-1	Platform	Install	04-AUG-2004 13:59:00
DEC AXPVMS TCP/IP V5.3-18	Full LP	Install	04-AUG-2004 13:59:00
DEC AXPVMS VMS V7.3-1	Oper System	Install	04-AUG-2004 13:59:00

```
8 items found
$
```

Installing the KGPSA-xx adapter card in an AlphaServer OpenVMS host

Complete this task to install the KGPSA-xx adapter card in a Hewlett-Packard AlphaServer OpenVMS host.

To install the adapter card, perform the following steps:

1. Shut down the AlphaServer OpenVMS host.
2. Use the procedures that are provided by the manufacturer of the adapter card to install the KGPSA-xx host adapter.
3. Bring the host system to a halt condition at the console level.
4. Type **set mode diag** to place the console in diagnostic mode.

Note: This step is only for AlphaServer models AS8x00, GS60/140, AS4x00, or AS1200. On all other machines, the **wwidmgr** command can be issued at any time prior to booting.

5. Type **wwidmgr -show adapter** to confirm that you installed each adapter properly.

Note: On AlphaServer models AS8x00, GS60/140, AS4x00, or AS1200, you need to set the console to diagnostic mode to use the **wwidmgr** command after initializing the system. On these systems, type **set mode diag** first as shown in step 4.

6. If necessary, update the adapter firmware.

The following is example output. This example shows the worldwide node name (WWN). You also need the worldwide port name (WWPN) to configure the storage unit host attachment. To determine the WWPN for the KGPSA adapters, replace "2" in the WWN with "1".

```
P00>>>set mode diag
Console is in diagnostic mode
P00>>>wwidmgr -show adapter
polling for units on kgpsa0, slot 9, bus 0, hose0...
kgpsaa0.0.0.9.0      PGA0      WWN 2000-0000-c922-69bf
polling for units on kgpsa1, slot 10, bus 0, hose0...
kgpsab0.0.0.10.0     PGB0      WWN 2000-0000-c921-df4b
item  adapter              WWN  Cur. Topo Next
Topo
[ 0] kgpsab0.0.0.10.0      2000-0000-c921-df4b    FABRIC
FABRIC
[ 1] kgpsaa0.0.0.9.0       2000-0000-c922-69bf    FABRIC
FABRIC
[9999] All of the above.
P00>>>
```

Setting the mode for the KGPSA-xx host adapter in an AlphaServer OpenVMS host

Complete this task to set the mode for the KGPSA-xx host adapter in a Hewlett-Packard AlphaServer OpenVMS host.

You must install the KGPSA-xx host adapter before you can set the mode.

To set the mode for the host adapter, perform the following steps:

1. Place the AlphaServer OpenVMS host into console mode.
2. The KGPSA-xx fibre-channel adapter card must be set to FABRIC. To verify that the Cur. Topo and Next Topo variables are set to FABRIC, type **wwidmgr -show adapter**.

The following is example output:

item	adapter	WWN	Cur. Topo	Next Topo
[0]	pga0.0.4.1	2000-0000-C923-1765	FABRIC	FABRIC

3. To set the KGPSA host adapter mode to FABRIC, type **wwidmgr -set adapter -item *item* -topo fabric**, where *item* is the value in brackets from the **wwidmgr -show adapter** command output.
4. To initialize the system, type **init**.
5. To verify the mode, type **wwidmgr -show adapter**.

Locating the WWPN for an OpenVMS host

This section describes how to locate the WWPN for a Hewlett-Packard AlphaServer or Integrity server OpenVMS host.

Locating the WWPN for an AlphaServer

Complete this task to locate the WWPN for a Hewlett-Packard AlphaServer OpenVMS host.

If the system is not running, you can find the WWPN by performing the following steps from the console:

1. From the console prompt, type **P0>>>wwidmgr -show ada**. The following is example output:

```
Probing timeout
item  adapter                WWN                Cur. Topo  Next Topo
[ 0] pga0.0.0.7.1          1000-0000-c922-d469    FABRIC     FABRIC
[ 1] pgb0.0.0.8.1          2000-0000-c922-6a63    FABRIC     FABRIC
[9999] All of the above.
```

You might receive the following error:

wwidmgr available only prior to booting. Reinit system and try again.

2. If you receive an error, type **P00>>>init** at the console prompt.

If the system is already running, you can find the WWPN by using the SHOW DEVICE command or the ANALYZE/SYSTEM utility. You must have the OpenVMS privilege CMKRNL to run the ANALYZE/SYSTEM utility.

- To use the SHOW DEVICE command, type **show device fg/full**. The following is example output. In this example, the WWPN is in the last line (1000-0000-C922-D469).

```
Device FGA0:, device type KGPSA Fibre Channel, is online, shareable, error
logging is enabled.

Error count          0      Operations completed          0
Owner process        ""      Owner UIC                  [SYSTEM]
Owner process ID     00000000  Dev Prot                   S:RWPL,0:RWPL,G,W
Reference count      0      Default buffer size        0
Current preferred CPU Id 0      Fastpath                   1
Current Interrupt CPU Id 0
FC Port Name 1000-0000-C922-D469  FC Node Name                2000-0000-C922-D469
```

- To use the ANALYZE/SYSTEM utility, perform the following steps:
 1. Type **ANALYZE/SYSTEM**.
 2. At the SDA> prompt, type **fc show dev fgadapter0**, where *adapter* is the letter that identifies the adapter; for example, **fc show dev fga0**. The following is example output. In this example, the WWPN is in the last line (1000-0000-C922-D469).

```
SDA> fc show dev fga0
FGA0: operational firmware revision SS3.20X7
port_name(adapter_id) = 1000-0000-C922-D469, node_name(host_id) = 2000-0000-C922-D469
```

Locating the WWPN for an Integrity server

Complete this task to locate the WWPN for a Hewlett-Packard Integrity server OpenVMS host.

If the system is not running, you can find the WWPN using the OpenVMS Extensible Firmware Interface (EFI) VMS_SHOW utility. For more information, see “Managing devices for an Integrity server using the OpenVMS EFI” on page 47.

If the system is already running, you can find the WWPN by using the SHOW DEVICE command or the ANALYZE/SYSTEM utility. You must have the OpenVMS privilege CMKRNL to run the ANALYZE/SYSTEM utility.

- To use the SHOW DEVICE command, type **show device fg/full**. The following is example output. In this example, the WWPN is in the last line (1000-0000-C922-D469).

```
Device FGA0:, device type KGPSA Fibre Channel, is online, shareable, error
logging is enabled.

Error count          0      Operations completed          0
Owner process        ""      Owner UIC                  [SYSTEM]
Owner process ID     00000000 Dev Prot                  S:RWPL,O:RWPL,G,W
Reference count      0      Default buffer size       0
Current preferred CPU Id 0      Fastpath                  1
Current Interrupt CPU Id 0
FC Port Name 1000-0000-C922-D469  FC Node Name      2000-0000-C922-D469
```

- To use the ANALYZE/SYSTEM utility, perform the following steps:
 1. Type **ANALYZE/SYSTEM**.
 2. At the SDA> prompt, type **fc show dev fgadapter0**, where *adapter* is the letter that identifies the adapter; for example, **fc show dev fga0**. The following is example output. In this example, the WWPN is in the last line (1000-0000-C922-D469).

```
SDA> fc show dev fga0
FGA0: operational firmware revision SS3.20X7
port_name(adapter_id) = 1000-0000-C922-D469, node_name(host_id) = 2000-0000-C922-D469
```

Managing devices for an Integrity server using the OpenVMS EFI

This section describes how you can use the OpenVMS Extensible Firmware Interface (EFI) to manage devices for the Hewlett-Packard Integrity server.

The following OpenVMS EFI utilities enable you to manage devices for an Integrity server from the EFI console, especially if you are using SAN boot:

VMS_SHOW

Displays all bootable devices that are mapped by the EFI console and their corresponding OpenVMS device names.

As shown in Figure 4 on page 48, the first output line shows the OpenVMS device name and additional information about the device. The additional information consists of vendor identification and product identification if the device is a disk, or a MAC address if the device is a network device. The second output line shows the file system designation (fsx) and its corresponding EFI device path.

```

VMS_SHOW.EFI dev
VMS: DQA0 IDE Drive
EFI: fs0: Acpi(HWP0002,0)/Pci(2|0)/Ata(Primary,Master)

VMS: EIA0 00-30-6E-F3-F2-52
EFI: Acpi(HWP0002,0)/Pci(3|0)/Mac(00306EF3F252)

VMS: DKA0 HP 36.4GST336753LC HPC4 V8.2-1
EFI: fs1: Acpi(HWP0002,100)/Pci(1|0)/Scsi(Pun0,Lun0)

VMS: EWA0 00-30-6E-F3-52-2C
EFI: Acpi(HWP0002,100)/Pci(2|0)/Mac(00306EF3522C)

VMS: DGA78 IBM 2107900 .531 V8.2-1
EFI: fs5: Acpi(HWP0002,300)/Pci(1|0)/Pci(4|0)/Fibre(WWN5005076303018020,Lun1000000000000)

VMS: DGA78 IBM 2107900 .531 V8.2-1
EFI: fs3: Acpi(HWP0002,300)/Pci(1|0)/Pci(4|0)/Fibre(WWN5005076303148020,Lun1000000000000)

```

Figure 4. OpenVMS Extensible Firmware Interface output

If specified, the utility matches the specified OpenVMS device name to the EFI console mapping. For multipath fibre-channel devices, the utility displays all paths that are associated with the specified OpenVMS device name.

There are several additional options:

- The `debug_dev` option displays the selected OpenVMS debug device.
- The `dump_dev` option displays the selected OpenVMS dump device for the dump-off-system-disk (DOSD) function.
- The `-fs` option displays the OpenVMS device names of those devices that have only the system disk.

VMS_BCFG

Adds an entry to the EFI Boot Manager using a specified OpenVMS device name.

The following is example output:

```

fs3:\efi\vms> VMS_BCFG.EFI boot show
The boot option list is:
01. VenHw(D65A6B8C-71E5-4DF0-A909-F0D2992B5AA9) "EFI Shell [Built-in]"
02. Acpi(HWP0002,300)/Pci(1|0)/Pci(4|0)/Fibre(WWN5005076303018020,Lun1000000000000)/HD(Part1,Sig0C516100-6657-11DC-AA2E-AA000400FEFF)/\efi\ms\ms_loader.efi "OpenVMS on $1DGA78: FGA0.5005-0763-0301-8020"
03. Acpi(HWP0002,300)/Pci(1|0)/Pci(4|0)/Fibre(WWN5005076303148020,Lun1000000000000)/HD(Part1,Sig0C516100-6657-11DC-AA2E-AA000400FEFF)/\efi\ms\ms_loader.efi "OpenVMS on $1DGA78: FGA0.5005-0763-0314-8020"
04. Acpi(HWP0002,100)/Pci(1|0)/Scsi(Pun0,Lun0)/HD(Part1,Sig76D23A51-9B8B-11DB-A618-AA000400FEFF)/\efi\ms\ms_loader.efi "DKA0 PKA0.0" OPT
05. Acpi(HWP0002,100)/Pci(1|0)/Scsi(Pun0,Lun0)/HD(Part1,Sig76D23A51-9B8B-11DB-A618-AA000400FEFF)/\efi\ms\ms_loader.efi "HP-UX Primary Boot: 0/1/1/0.0.0"
06. Acpi(HWP0002,0)/Pci(2|0)/Ata(Primary,Master) "CDROM" OPT

fs3:\efi\ms>

```

Confirming fibre-channel switch connectivity for OpenVMS

Complete this task to confirm fibre-channel switch connectivity for the Hewlett-Packard OpenVMS host system.

1. Open a telnet session and log in to the switch as an administrator.
2. To confirm that each storage unit host adapter has performed a fabric login to the switch, type **switchshow**. The following is example output:

```

snj2109f16h4:os1> switchshow
switchName:      snj2109f16h4
switchType:      9.1
switchState:     Online
switchRole:      Principal
switchDomain:     1
switchId:        fffc01
switchWwn:       10:00:00:60:69:50:0c:3e
switchBeacon:    OFF
port 0: id N1 Online      F-Port 50:05:07:63:00:c9:91:62
port 1: id N1 Online      F-Port 10:00:00:00:c9:22:d2:08
port 2: id N1 Online      F-Port 10:00:00:00:c9:22:6a:63
port 3: id N1 Online      F-Port 50:00:1f:e1:00:00:2b:11
port 4: id N2 No_Light
port 5: id N1 Online      F-Port 10:00:00:00:c9:22:d4:69
port 6: id N1 Online      F-Port 10:00:00:00:c9:22:67:38
port 7: id N1 Online      L-Port 1 private, 3 phantom
port 8: id N2 No_Light
port 9: id N1 Online      F-Port 10:00:00:00:c9:22:69:bf
port 10: id N1 Online     F-Port 10:00:00:00:c9:21:df:4b
port 11: id N1 Online     F-Port 50:05:07:63:00:cf:8d:7e
port 12: id N2 No_Light
port 13: id N1 Online     F-Port 50:05:07:63:00:c7:91:1c
port 14: id N1 Online     F-Port 50:05:07:63:00:cd:91:62
port 15: -- N2 No_Module
snj2109f16h4:os1>

```

Defining a UDID for OpenVMS storage unit volumes

Each OpenVMS fibre-attached volume requires a user-defined identifier or unit device identifier (UDID). This section describes how you define a UDID for each storage unit volume.

A UDID is a nonnegative integer that is used in the creation of the OpenVMS device name. All fibre-attached volumes have an allocation class of \$1\$, followed by the letters, DGA, followed by the UDID. All storage unit LUNs that you assign to an OpenVMS system need an UDID so that the operating system can detect and name the device. LUN 0 must also have a UDID; however, the system displays LUN 0 as \$1\$GGA<UDID>, not as \$1\$DGA<UDID>. See the HP document, *Guidelines for OpenVMS Cluster Configurations* for more information about fibre-attached storage devices.

You can use the IBM System Storage DS Storage Manager or the DS CLI to set a value in a storage unit volume name field that is used by OpenVMS systems as the UDID for the volume. DS CLI examples are provided in this section. You can find the DS CLI on the CD that you receive with the storage unit. See the *IBM System Storage DS6000 Command-Line Interface User's Guide* for more information.

The DS CLI is a general purpose utility that supports various storage unit functions. The DS CLI allows 16 alphanumeric characters as input when you complete the storage unit volume name field. OpenVMS UDID values must be an integer within the range of 0 to 32767. Therefore, you must ensure that the input is valid for UDID support. The utility does not enforce UDID rules. It accepts values, such as AaBbCcDd, that are not valid for OpenVMS. It is possible to assign the same UDID value to multiple storage unit volumes. However, each volume that you assign to an OpenVMS system must have a value that is unique for that system or throughout the OpenVMS cluster to which that system belongs. Review the HP OpenVMS documentation for UDID rules to verify that your input is valid.

Note: Volumes with UDIDs greater than 9999 cannot be MSCP-served in an OpenVMS cluster to other systems.

Example

The following example uses the DS CLI to add or change a name to an existing DS volume. In this example, the DS CLI is in interactive mode and a configuration profile file has been defined. The final command uses the AlphaServer console to list fibre-attached volumes.

1. Use the `chfbvol` command to change the name of a fixed block volume. For example, to set the UIDD value to 21, type **`chfbvol -name 21 0001`**.

The value for the name parameter in the DS CLI command is the UDID field for the AlphaServer. This command returns the volume ID. The following is example output:

```
CMUC00026I chfbvol: FB volume 0001 successfully modified.
```

Note: The first volume, LUN 0, will be reported as a CCL device and not as a disk volume.

2. To make a volume group called `VMS_A0` and add a volume to it, type the following command:

```
mkvolgrp -type scsimap256 -volume 0001 VMS_A0
```

This command returns the volume group ID. The following is example output:

```
CMUC00030I mkvolgrp: Volume group V0 successfully created.
```

3. To create an OpenVMS host with the DS CLI and associate a volume group with it, type the following command:

```
mkhostconnect -wwname 10000000ABCDEF98 -hosttype HpVms -volgrp v0 ES40_A
```

This command returns the host connection ID. The following is example output:

```
CMUC00012I mkhostconnect: Host connection 0005 successfully created.
```

4. To display the defined attributes for a host connection, type **`showhostconnect 0005`**.

The following is example output:

```
Name      ES40_A
ID        0005
WWPN      10000000ABCDEF98
HostType   HpVms
LBS       512
addrDiscovery LUNPolling
Profile    HP - Open VMS
portgrp    0 volgrpID
V0 atchtopo -
ESSIOport  all
```

5. To display the volumes in a volume group and its attributes, type **`showvolgrp v0`**.

The following is example output:


```
Name VMS_A0
ID V0
Type SCSI Map 256
Vols 002A 0000F 0001
```

6. Use the `wwidmgr -show wwid` command at the AlphaServer console to list fibre-attached volumes that have been detected by its fibre-channel host adapters. If a volume has no UDID or has a UDID that is not valid, the volume UDID is minus one (-1). When it is booted, OpenVMS does not detect a volume with -1 as a UDID number. For more detail, see “Confirming fibre-channel switch connectivity for OpenVMS” on page 48.

Accessing the storage unit volumes from an OpenVMS host

Complete this task to access the storage unit volumes from a Hewlett-Packard OpenVMS host.

To access the storage unit volumes, perform the following steps:

1. Start the host operating system.

During the startup procedure, the operating system scans for fibre-channel storage devices and configures them automatically. You can display the fibre-channel disks with the following `SHOW DEVICE` command:

```
$ SHOW DEVICE DGA
```

If you have assigned new storage unit volumes to the host when the operating system was already running, or if the `SHOW DEVICE` command does not show newly-assigned DS6000 disks, you can use the `SYSMAN` utility to configure the new devices manually:

```
RUN SYS$SYSTEM:SYSMAN
SYSMAN> IO AUTOCONFIGURE/LOG
SYSMAN> EXIT
$ SHOW DEVICE DGA
```

2. Initialize the storage unit volumes.

On the OpenVMS platform, you can initialize storage unit volumes as ODS-2 or ODS-5 volumes. You can use the volumes to create volume sets. Volume sets are concatenated volumes that form a larger volume. See the HP document, *OpenVMS System Manager's Manual, Volume 1: Essentials*. For OpenVMS shadow sets, HP recommends that you add the qualifiers `/SHADOW/ERASE` to the `INITIALIZE` command. For more information, see the HP document, *Volume Shadowing for OpenVMS*.

If you plan to use DS6000 Logical Volume Expansion, you can prepare the volume with the one-time allocation of extra bitmap space and configure the maximum size that will ever be used on this volume. You can perform this task at disk initialization time with the `INITIALIZE/LIMIT` command or on a privately-mounted volume with the `SET VOLUME/LIMIT` command. These commands require OpenVMS Alpha version 7.3-2 or later. For more information, see the HP document, *Volume Shadowing for OpenVMS*.

3. Mount the storage unit volumes.

For OpenVMS shadow sets, the storage unit does not support `READL` or `WRITEL` commands. Therefore, the volume does not support the shadowing data repair (disk bad block errors) capability as some other disks do. Add the `/OVERRIDE=NO_FORCED_ERROR` qualifier to the `MOUNT` command when you use storage unit volumes as a shadow set. This qualifier suppresses bad block handling by OpenVMS shadowing data repair. For more information, see the HP document, *Volume Shadowing for OpenVMS*.

4. Access the storage unit volumes.

The following is an example showing what is displayed when you use the standard OpenVMS storage configuration utilities (SHOW DEVICES DG command).

```
$ SHOW DEVICES DG

Device      Device Error Volume Free   Trans Mnt
Name        Status Count Label  Blocks Count Cnt
1$DGA20: (HANK) Online    0
1$DGA21: (HANK) Online    0
$ INITIALIZE/SYSTEM 1$DGA20 ESS001
$ MOUNT/SYSTEM 1$DGA20 ESS001
%MOUNT-I-MOUNTED, VOLUME_ESS001 mounted on 1$DGA20: (HANK)
$ DIRECTORY 1$DGA20:[000000]

000000.DIR;1      BACKUP.SYS;1      BADBLK.SYS;1      BADLOG.SYS;1
BITMAP.SYS;1      CONTIN.SYS;1      CORIMG.SYS;1      INDEXF.SYS;1
SECURITY.SYS;1    VOLSET.SYS;1

Total of 10 files.
$ DISMOUNT 1$DGA20
$
```

Guidelines for OpenVMS LUN 0

This section describes guidelines for LUN 0 on an OpenVMS host system. OpenVMS considers the storage unit volume that becomes LUN 0 as the Command Console LUN (CCL), or pass-through LUN.

The storage unit assigns LUN numbers using the lowest available number. The first storage unit volume that is assigned to a host is LUN 0, the next volume is LUN 1, and so on. When a storage unit volume is unassigned, the system reuses the LUN number that it had when assigning another volume to that host.

You can verify the LUN number assignment with the DS CLI `showvolgrp -lunmap` command and modify the LUN number assignment with the `chvolgrp -lun` command.

The storage unit does not support CCL command functions. This storage unit volume (LUN 0 for OpenVMS) does not display when you issue the `wwidmgr -show wwid` AlphaServer console command. When OpenVMS is running and a UDID has been set for that storage unit volume, the storage unit volume LUN 0 displays as a GGA device type, not as a DGA device. Although OpenVMS does not strictly require a UDID for the CCL, the `SHOW DEVICE GGA` command displays CCL devices if you set a UDID. You can display the multiple paths and diagnose failed paths to the storage controller using the `SHOW DEVICE/FULL GGA` command.

Guidelines:

The first storage unit volume that you assign to an OpenVMS system becomes LUN 0. However, the volume that you assign can be used by the system for only support capability and should be at the minimum size. An OpenVMS host cannot use the volume for any other purpose. Multiple OpenVMS hosts, even in different clusters that access the same storage unit, can share the same storage unit volume as LUN 0, because there will be no other activity to this volume.

When you issue the `wwidmgr -show wwid` command at the AlphaServer console, LUN 0 does not display. Only fibre-channel storage devices are listed in the

output. The storage unit LUN 0 is presented as a CCL device, and therefore, is not shown. You can only see the LUN 0 from the operating system when a UDID had been assigned.

Restrictions for OpenVMS fibre-channel connections

The restrictions listed in this section apply only if your DS6000 licensed machine code has a version before 5.0.3.157.

Note: You can display the versions of the DS CLI, the Storage Manager, and the licensed machine code by using the DS CLI `ver -l` command.

The following restrictions are required for the DS6000 host adapter to maintain compatibility with the OpenVMS host system. Compatibility is enabled on a DS6000 adapter port after a defined host establishes fibre-channel connectivity to one or more storage unit volumes.

- You must dedicate DS6000 adapter ports for only the OpenVMS type. It is recommended that each OpenVMS host adapter be in a fabric zone with one or more DS6000 adapter ports.
- All DS6000 adapter ports that are in a fabric zone with an OpenVMS host adapter must have at least one DS6000 volume assigned to it. This can be the LUN 0 volume.
- Multiple OpenVMS systems can access the same DS6000 adapter port. However, you must define each system for that specific DS6000 port and assign at least one DS6000 volume.
- To re-enable compatibility, you can force the defined OpenVMS host to reestablish connectivity to the DS6000 adapter by disabling the switch port it is connected to, and then enabling the switch port again.

The following problems are a few examples of issues that might occur when you attach a storage unit to an HP AlphaServer OpenVMS host system without fulfilling the OpenVMS fibre-channel restrictions.

If the system hangs during boot

Your system might be attempting to access a host adapter port that is not configured for it. To correct your system, check to ensure that the following items are true:

1. The fibre-channel switch is enabled to ensure zoning.
2. The zones for the OpenVMS host are connected to only those ports that are correctly configured to support them.
3. Other fabric zones do not include the OpenVMS host adapters by mistake.

If the system reaches the mount verification timeout

All pending and future I/O requests to the volume fail. You have to dismount and remount the disk before you can access it again.

The mount verification process for fibre attached volumes might not complete until the system reaches the mount verification timeout. This can be caused by the same scenario that causes the system to hang during boot. Verify that the path to the affected volumes is available, and then follow the above list.

If another host system gains access to a storage unit host adapter port that is dedicated for OpenVMS

The port will have its compatibility mode disabled. This can occur when the fabric switch has disabled all zoning. After you reenable zoning, the storage unit port compatibility remains disabled. Some internal storage unit processes cause a momentary loss of connectivity to volumes with activity I/O. OpenVMS has no problems when the storage unit adapter port is in compatibility mode. If the mode is disabled, disk I/O could fail a read or write and display the following message:

-SYSTEM-F-TOOMANYRED, too many redirects

Force the host adapter to reestablish connectivity

Forcing the OpenVMS host adapter to reestablish fibre-channel connectivity with the storage unit adapter port will enable compatibility mode. You can force the connection only by disconnecting and then reconnecting one end of the physical fibre-channel cable between the host and the affected storage unit adapter port. You can also reestablish connectivity by accessing the fibre-channel switch and disabling one of the switch ports, and then enabling it again.

HP AlphaServer Tru64 UNIX host attachment

This chapter describes how you attach a Hewlett-Packard (HP) AlphaServer Tru64 UNIX[®] host system to a DS6000 with fibre-channel adapters.

This chapter contains the following sections:

- Confirming the installation of the Tru64 UNIX operating system
- Installing and setting the mode for the KGPSA-xx adapter card in a Tru64 UNIX host system
- Locating the WWPN for a Tru64 UNIX host
- Configuring a Tru64 UNIX host
- Configuring storage for a fibre-channel Tru64 UNIX host
- Removing persistent reserves for Tru64 UNIX 5.x
- Limitations for Tru64 UNIX

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, and adapters
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Notes:

1. You do not need the Subsystem Device Driver because Tru64 UNIX manages multipathing.
2. IBM supports cascaded switches in configurations up to a maximum of 8 switches with a maximum of 3 interswitch hops for any path. Use two hops for normal operations with the third hop reserved for backup paths.
3. The DS CLI is supported on Tru64 UNIX.

Confirming the installation of the Tru64 UNIX operating system

Complete this task to confirm the installation of the Hewlett-Packard Tru64 UNIX operating system.

If you use the storage unit volumes member boot disks for a clustered or nonclustered configuration, install the operating system from the console level. You can use the storage unit LUN as a boot disk only for the Tru64 5.x operating system.

1. Confirm the installation of the appropriate version of Tru64 UNIX. For the Tru64 UNIX 5.x operating system, type **sizer -v** to confirm installation.

The following is command output:

```
# sizer -v
v5.x
#
Compaq Tru64 UNIX V5.1A (Rev. 1885); Tue Sept 24 14:12:40 PDT 2002
```

2. Ensure that you are at patch level 7 and that all kernel options are active. For Tru64 UNIX cluster configurations, you must install patch 399.00 with patch security (SSRT0700U). For information about other patches for this version, refer to the documentation from HP.

Installing the KGPSA-xx adapter card in a Tru64 UNIX host system

Complete this task to install the KGPSA-xx adapter card in a Hewlett-Packard Tru64 UNIX host system.

To install the adapter card, perform the following steps:

1. Shut down the HP AlphaServer host system.
2. Use the procedures provided by the manufacturer of the adapter card to install the KGPSA-xx host adapter.
3. Restart the host (nonclustered configurations) or each cluster member (clustered configurations).
4. Bring each host system to a halt condition at the console level.
5. Type **set mode diag** at the HP AlphaServer console to place the console in diagnostic mode.
6. Type **wwidmgr-show adapter** to confirm that you installed each adapter properly.
7. If necessary, update the adapter firmware.

The following is example output. This example shows the worldwide node name (WWNN). You need the worldwide port name (WWPN) to configure the storage unit host attachment. To determine the WWPN for the KGPSA\FCA adapters, replace the "2" in the WWNN with a "1".

```
P00>>>set mode diag
Console is in diagnostic mode
P00>>>wwidmgr -show adapter
polling for units on kgpsa0, slot 9, bus 0, hose0...
kgpsaa0.0.0.9.0      PGA0      WWN 2000-0000-c922-69bf
polling for units on kgpsa1, slot 10, bus 0, hose0...
kgpsab0.0.0.10.0     PGB0      WWN 2000-0000-c921-df4b
item  adapter          WWN  Cur. Topo Next
Topo
[ 0] kgpsab0.0.0.10.0    2000-0000-c921-df4b    FABRIC
FABRIC
[ 1] kgpsaa0.0.0.9.0     2000-0000-c922-69bf    FABRIC
FABRIC
[9999] All of the above.
P00>>>
```

Setting the mode for the KGPSA-xx host adapter in a Tru64 UNIX host system

Complete this task to set the mode for the KGPSA-xx host adapter in a Hewlett-Packard Tru64 UNIX host system.

You must install the KGPSA-xx host adapter before you can set the mode.

The default KGPSA/FCA mode setting is FABRIC, so that directly attaching the AlphaServer to the storage unit using fibre-channel KGPSA-xx adapters does not work without modification. You must change the mode setting to LOOP mode.

To set the mode for the host adapter, perform the following steps:

1. Type `# shutdown -h now` to shut down the operating system.
2. Type `init` to initialize the system.
3. Type `wwidmgr -show adapter` to check the mode. The following is example output:

```
item  adapter      WWN      Cur. Topo  Next Topo
[ 0]  pga0.0.4.1    2000-0000-C923-1765  FABRIC    FABRIC
```

4. Type one of the following commands to set the mode of the KGPSA/FCA host adapter:
 - For FABRIC mode, type `wwidmgr -set adapter -item 9999 -topo fabric`
 - For LOOP mode, type `wwidmgr -set adapter -item 9999 -topo loop`
5. Type `init` to initialize the system.
6. Type `wwidmgr -show adapter` to check the mode.
7. Use the IBM System Storage DS Storage Manager to set your port attributes to match your host settings:
 - **For arbitrated loop:** Set your port attribute to Direct Connect.
 - **For point-to-point:** Set your port attribute to Switched Fabric.

Locating the WWPN for a Tru64 UNIX host

Complete this task to locate the WWPN for a Hewlett-Packard AlphaServer Tru64 UNIX host.

1. From the console prompt, type **P0>>>wwidmgr -show ada**.
The following is example output:

```
Probing timeout
item  adapter                WWN                Cur. Topo  Next Topo
[ 0]  pga0.0.0.7.1          1000-0000-c922-d469  FABRIC     FABRIC
[ 1]  pgb0.0.0.8.1          2000-0000-c922-6a63  FABRIC     FABRIC
[9999] All of the above.
```

You might receive the following errors:

Message:

wwidmgr available only prior to booting. Reinit system and try again.

Action:

Type **P00>>>init** at the prompt.

If the system is already running, you can find the WWPN in the log file
/var/adm/messages.

2. Type **#fgrep wwn /var/adm/messages**.

The following is example output. You can find the WWPN in the last column.

```
...
Nov  9 09:01:16 osplcpq-ds20 vmunix: KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
Nov 10 10:07:12 osplcpq-ds20 vmunix: KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
Nov 13 17:25:28 osplcpq-ds20 vmunix: KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
Nov 14 11:08:16 osplcpq-ds20 vmunix: KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
Nov 15 10:49:31 osplcpq-ds20 vmunix: KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
...
```

Configuring a Tru64 UNIX host

The following sections describe how you configure a Hewlett-Packard AlphaServer Tru64 UNIX host with fibre-channel adapters.

Confirming fibre-channel switch connectivity for Tru64 UNIX

Complete this task to confirm fibre-channel switch connectivity for a Hewlett-Packard AlphaServer Tru64 UNIX host.

To confirm fibre-channel switch connectivity, perform the following steps:

1. Open a telnet session and log in to the switch as an administrator.
2. Confirm that each host adapter has performed a fabric login to the switch.
3. Type switchshow to confirm that each storage unit host adapter has performed a fabric login to the switch. The following is example output:


```

snj2109f16h4:os1> switchshow
switchName:      snj2109f16h4
switchType:      9.1
switchState:     Online
switchRole:      Principal
switchDomain:     1
switchId:        fffc01
switchWwn:       10:00:00:60:69:50:0c:3e
switchBeacon:    OFF
port 0: id N1 Online      F-Port 50:05:07:63:00:c9:91:62
port 1: id N1 Online      F-Port 10:00:00:00:c9:22:d2:08
port 2: id N1 Online      F-Port 10:00:00:00:c9:22:6a:63
port 3: id N1 Online      F-Port 50:00:1f:e1:00:00:2b:11
port 4: id N2 No_Light
port 5: id N1 Online      F-Port 10:00:00:00:c9:22:d4:69
port 6: id N1 Online      F-Port 10:00:00:00:c9:22:67:38
port 7: id N1 Online      L-Port 1 private, 3 phantom
port 8: id N2 No_Light
port 9: id N1 Online      F-Port 10:00:00:00:c9:22:69:bf
port 10: id N1 Online     F-Port 10:00:00:00:c9:21:df:4b
port 11: id N1 Online     F-Port 50:05:07:63:00:cf:8d:7e
port 12: id N2 No_Light
port 13: id N1 Online     F-Port 50:05:07:63:00:c7:91:1c
port 14: id N1 Online     F-Port 50:05:07:63:00:cd:91:62
port 15: -- N2 No_Module
snj2109f16h4:os1>

```

Confirming fibre-channel storage connectivity for Tru64 UNIX

Complete this task to confirm fibre-channel storage connectivity for a Hewlett-Packard AlphaServer Tru64 UNIX host.

To confirm fibre-channel storage connectivity, perform the following steps:

1. Reset the host (nonclustered configurations) or each cluster member (clustered configurations).
2. Bring each host system to a halt condition at the console level.
3. Type set mode diag at the HP AlphaServer console (if required by the host) to place the console in diagnostic mode.

Type wwidmgr -show wwid to display the information about the storage unit volume at the console level. You can use this information to identify the volumes that are attached to an AlphaServer. The following example output shows information about the storage unit volumes that you can see at the AlphaServer console.

```

P00>>>set mode diag
Console is in diagnostic mode
P00>>>wwidmgr -show wwid
[0] UDID: -1 WWID:01000010:6000-1fe1-4942-4d20-0000-0000-28b1-5660 (ev:none)
[1] UDID: -1 WWID:01000010:6000-1fe1-4942-4d20-0000-0000-2881-5660 (ev:none)
[2] UDID: -1 WWID:01000010:6000-1fe1-4942-4d20-0000-0000-2821-5660 (ev:none)
P00>>>

```

4. Type wwidmgr -show adapter to confirm storage attachment.

See “Tru64 UNIX UDID hexadecimal representations” for an explanation of the UDID.

Tru64 UNIX UDID hexadecimal representations

This section describes UDID representations for storage unit volumes.

The UDID for each volume appears as -1, signifying that the UDID is undefined. With the supported storage unit LMC, all UDIDs for storage unit volumes are undefined.

The underscore in the following example highlights the hex string that identifies a storage unit volume that is attached to an AlphaServer.

```
01000010:6000-1fe1-4942-4d20-0000-0000-2821-5660
```

The third and fourth quartet of the UDID number is always the value “4942-4d20”. This is the string IBMb in hex and represents a storage unit volume.

The underscore in the following example highlights the hex string that identifies the decimal volume number of the storage unit volume. The first three characters of the next to last quartet of numbers is the hex string representation. The following example shows that the storage unit volume number is decimal 282.

```
01000010:6000-1fe1-4942-4d20-0000-0000-2821-5660
```

The following example shows a hex representation of the last 5 characters of the storage unit volume serial number.

```
01000010:6000-1fe1-4942-4d20-0000-0000-2821-5660
```

Preparing to boot from the storage unit for a Tru64 UNIX host

Complete this task in preparation for booting from the storage unit for a Hewlett-Packard AlphaServer Tru64 UNIX host system.

Use the **wwidmgr** command to set up each device that you use for booting or dumping. After you set up a device, the console retains the information that is needed to access the device in nonvolatile memory.

Rerun the **wwidmgr** command if the system configuration changes and the nonvolatile information is no longer valid.

1. Display the WWIDs of all assigned storage unit volumes with the **wwidmgr -show wwid** command.
2. Determine which storage unit volume that you want to use as a boot or dump device by decoding the serial number as described in “Tru64 UNIX UDID hexadecimal representations” on page 58.
3. Assign a unit number with the **wwidmgr -quickset -item i -unit u** command, where *i* is the wwidmgr item number and *u* is the unit number you choose. You can find the item number inside the square brackets of the output from the **wwidmgr -show wwid** command.
4. Reinitialize the server with the **init** command.
When you make changes with the **wwidmgr** command, they do not take effect until the next system initialization.
5. Use the **show device** command to verify that the system displays the disk as console device **DGAu**, with the unit number that you defined.

After the initialization, the console **show device** command displays each fibre-channel adapter followed by the paths through that adapter to each of the defined fibre-channel disks. The path-independent device name for each fibre-channel disk is displayed in the second column.

- Set the default boot device console variable, **bootdef_dev**, to match the console device name of the boot disk. In a multipath environment, use the asterisk (*) as a wildcard to make all boot paths available. An example of the multipath command is **set bootdef_dev DGAu.***, where *u* is the unit number and * denotes all possible paths.

Configuring kernel SCSI parameters on a Tru64 UNIX host

Complete this task to set up and configure the Hewlett-Packard AlphaServer Tru64 UNIX device parameter database.

This procedure applies to all versions of Tru64 UNIX 4.0f and later. See the Tru64 UNIX man files for **ddr.dbase** and **ddr_config** for more information.

- With the storage quiesced, the host system in single-user mode, and logged in to the host system with root privileges, edit the **/etc/ddr.dbase** file, and include the following parameter values in the **DISKS** subsection.

Table 4 provides the recommended settings for the DS6000 SCSI device.

Table 4. Recommended Tru64 UNIX settings for the SCSI DEVICE.

Parameters	Recommended settings
TypeSubClass	hard_disk, raid
BadBlockRecovery	disabled
DynamicGeometry	true
LongTimeoutRetry	enabled
TagQueueDepth	20
ReadyTimeSeconds	180
CMD_Write_Verify	supported
InquiryLength	255
RequestSenseLength	255

- Compile by executing the following command: **ddr_config -c** .
- Confirm that the values were included by issuing the following command: **ddr_config -s disk IBM 1750** .
- Edit the **/sys/data/cam_data.c** file to change the nonread/write command timeout values in the changeable disk driver timeout section from 10 seconds to 60. Locate the existing timeout value **u_long cdisk_to_def = 10; /10* seconds */**, and change the value to **u_long cdisk_to_def = 60; /60* seconds */**.
- Compile **cam_data.c** by issuing the following command: **doconfig -c HOSTNAME**, where HOSTNAME is the name of the system kernel. It can be found in the **/sys/conf/** directory.

The following is example output:

```
#doconfig -c ES47
*** KERNEL CONFIGURATION AND BUILD PROCEDURE ***

Saving /sys/conf/ES47 as /sys/conf/ES47.bck

Do you want to edit the configuration file? (y/n) [n]: y

Using ed to edit the configuration file. Press return when ready,
or type 'quit' to skip the editing session: quit

*** PERFORMING KERNEL BUILD ***
Working....Wed Mar 22 17:36:19 PST 2006

The new kernel is /sys/ES47/vmunix
#
```

Verifying fibre-channel attachment to a Tru64 UNIX host

Complete this task to verify the fibre-channel attachment of the storage unit volumes to a Hewlett-Packard AlphaServer Tru64 UNIX host system.

To verify the fibre-channel attachment to a Tru64 UNIX 5.x host, perform the following steps:

1. Use the **hwmgr** command to verify the attachment of the storage unit volumes for Tru64 5.x. The following example shows the commands you can use to verify the attachment of the storage unit volumes.

```
# hwmgr -view dev -cat disk
HWID: Device Name      Mfg      Model      Location
-----
54: /dev/disk/floppy0c    3.55in floppy  fdi0-unit-0
60: /dev/disk/dsk1c  DEC      RZ2DD-LS (C) DEC  bus-2-targ-0-lun-0
63: /dev/disk/cdrom0c  COMPAQ    CDR-8435      bus-5-targ-0-lun-0
66: /dev/disk/dsk5c  IBM      2105F20      bus-0-targ-253-lun-0
67: /dev/disk/dsk6c  IBM      2105F20      bus-0-targ-253-lun-1
68: /dev/disk/dsk7c  IBM      2105F20      bus-0-targ-253-lun-2
:
:
# hwmgr -get attributes -id 66
66:
name = SCSI-WWID:01000010:6000-1fe1-0000-2b10-0009-9010-0323-0046
category = disk
sub_category = generic
architecture = SCSI
:
:
```

2. Use the example Korn shell script `dsvol`, which is shown in the following example, to display a summary that includes information for all the storage unit volumes that are attached.

```
echo Extracting DS volume information...
for ID in `hwmgr -view dev -cat disk | grep ibm1750 | awk '{ print $1}'`
do echo; echo DS vol, H/W ID $ID
hwmgr -get attrib -id $ID | awk '/phys_loc//dev_base//capacity//serial/'
done
```

Note: In the following example output, you can see storage unit volumes 282, 283, and 284 as LUNS 0, 1, and 2 respectively. You can access the LUNs in the Tru64 UNIX host by using the following special device files:

- /dev/rdisk/dsk3
- /dev/rdisk/dsk4
- /dev/rdisk/dsk5

```
# ./dsvol | more
Extracting DS volume information...ESS vol, H/W ID 38:
phys_location = bus-2-targ-0-lun-0
dev_base_name = dsk3
capacity = 5859392
serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-2821-5660

DS vol, H/W ID 39:
phys_location = bus-2-targ-0-lun-1
dev_base_name = dsk4
capacity = 5859392
serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-2831-5660

DS vol, H/W ID 40:
phys_location = bus-2-targ-0-lun-2
dev_base_name = dsk5
capacity = 5859392
serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-2841-5660
#
```

Configuring storage for a fibre-channel Tru64 UNIX host

Complete this task to configure the storage for a fibre-channel Hewlett-Packard AlphaServer Tru64 UNIX host.

You can use the standard Tru64 UNIX storage configuration utilities to partition and prepare storage unit LUNs, and to create and mount file systems.

Perform one of the following sets of steps to configure the storage for a Tru64 UNIX 5.x file system.

- To configure an AdvFS file system:
 1. Type: # disklabel -wr /dev/rdisk/dsk6c
 2. Type: # mkfdmn /dev/disk/dsk6c adomain
 3. Type: # mkfset adomain afs
 4. Type: # mkdir /fs
 5. Type: # mount -t advfs adomain#afs /fs
- To configure a Ufs file system:
 1. Type: # disklabel -wr /dev/disk/dsk6c
 2. Type: # newfs /dev/disk/dsk6c
 3. Type: # mkdir /fs
 4. Type: # mount -t ufs /dev/disk/dsk6c /fs

Removing persistent reserves for Tru64 UNIX 5.x

Complete this task to remove persistent reserves for a Hewlett-Packard Tru64 UNIX 5.x host system.

In the clustered environment, Tru64 5.x hosts place a persistent reserve whenever you assign a LUN. If you perform a FlashCopy or Remote Mirror and Copy

operation on a LUN that has persistent reserve, it will fail. If you want to perform a FlashCopy or Remote Mirror and Copy operation, remove the persistent reserve on the target LUN before you assign the LUN.

For example, assume that there are two Tru64 5.x hosts called Alpha1 and Alpha2 that can access storage unit LUNs with fibre connections. One fibre connection goes from the KGPSA-xx in Alpha1 to the switch. One fibre connection goes from Alpha2 to the switch. Another fibre connection goes from the switch to the storage unit.

Use a storage unit volume, for example 10a-21380, as a target LUN to perform a FlashCopy. There are two hosts and four connections: two from Alpha1 and two from Alpha2. The storage unit volume 10a-21380 has four registrants. One registrant is reserved. Use the `essvol` script to find the devices that are associated with storage unit volume 10a-21380 on each Trucluster node.

The following example shows how to remove a persistent reserve using the `essvol` script:

```
alpha1> essvol
DS vol, H/W ID 176:
  phys_location = bus-9-targ-0-lun-0
  dev_base_name = dsk43
  capacity = 3906304
  serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-1042-1380

DS vol, H/W ID 225:
  phys_location = bus-9-targ-0-lun-7
  dev_base_name = dsk47
  capacity = 3906304
  serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-10a2-1380

alpha2> essvol
DS vol, H/W ID 176:
  phys_location = bus-4-targ-0-lun-0
  dev_base_name = dsk43

capacity = 3906304
serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-1042-1380

DS vol, H/W ID 225:
  phys_location = bus-3-targ-0-lun-1
  dev_base_name = dsk47
  capacity = 3906304
  serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-10a2-1380
```

Use the `scu` command to see the reservations on these devices. The following examples show what you see when you use the `scu` command. You can associate `dsk47` on Alpha1 and Alpha2 with storage unit volume 10a-21380.

```
alpha1> scu -f /dev/rdisk/dsk47c show reservations
Persistent Reservation Header:
```

```
    Generation Value: 49
    Additional Length: 16
```

```
Reservation Descriptors:
```

```
    Reservation Key: 0x30001
    Scope-Specific Address: 0
    Reservation Type: 0x5 (Write Exclusive Registrants Only)
    Reservation Scope: 0 (LU - full logical unit)
    Extent Length: 0
```

```
alpha2> scu -f /dev/rdisk/dsk47c show reservations
Persistent Reservation Header:
```

```
    Generation Value: 49
    Additional Length: 16
```

```
Reservation Descriptors:
```

```
    Reservation Key: 0x30001
    Scope-Specific Address: 0
    Reservation Type: 0x5 (Write Exclusive Registrants Only)
    Reservation Scope: 0 (LU - full logical unit)
    Extent Length: 0
```

Each device shows 0x30001 as a reservation key. You must issue the **scu** command again to remove the persistent reserve using reservation key 0x30001 on each node. Tru64 places a unique reservation key on each LUN whenever the storage unit assigns the LUN. The reservation key can only be removed from that specific host to which it is assigned. Because it is not possible to tell exactly which registrant on the host holds the reservation, you must issue an **scu clear** command on each node in the cluster. The following is example output. One of the two commands that you see in this example clears the persistent reserve on storage unit volume 10a-21380.

```
alpha1> scu -f /dev/rdisk/dsk47c press clear key 0x30001
alpha2> scu -f /dev/rdisk/dsk47c press clear key 0x30001
```

Use the **scu** command to check reservations. The following is example output. One of the two commands you see in this example clears the persistent reserves on storage unit volume 10a-21380.

```
alpha1> scu -f /dev/rdisk/dsk47c show reservations
```

```
Persistent Reservation Header:
```

```
Generation Value: 50
```

```
Additional Length: 0
```

```
Reservation Descriptors:
```

```
Alpha2> scu -f /dev/rdisk/dsk47c show reservations
```

```
Persistent Reservation Header:
```

```
Generation Value: 50
```

```
Additional Length: 0
```

```
Reservation Descriptors:
```

After removing the persistent reserve from a storage unit volume, you can use it as a target LUN for FlashCopy or Remote Mirror and Copy.

Limitations for Tru64 UNIX

This section lists the limitations for Tru64 UNIX for fibre-channel connections.

Boot Volumes

IBM does not support FlashCopy or Remote Mirror and Copy on Tru64 boot volumes (cluster boot volumes). Do not attempt to clear persistent reserves on these LUNs.

UFS file system

The data will be inconsistent if you perform a FlashCopy on a target LUN which is online. Take the LUN offline before you perform a FlashCopy.

AdvFS file system

It is not possible to access a FlashCopy target volume on the same host as the source because of domain/fileset advfs concepts. You must unmount the source volume before you can access the target volume.

From the command line, type the following commands:

1. # umount /source
2. # mkdir /etc/fdmns/t_domain (target domain)
3. # ln -s /dev/disk/dsk47c dsk47c (target volume)
4. # mkdir /target
5. # mount -t advfs dsk47c
6. # source /target (source is fileset of source volume)

Configuring AdvFS parameters on a Tru64 UNIX host

Complete this task to configure the Advance File System (AdvFS) parameters on a Hewlett-Packard AlphaServer Tru64 UNIX host.

Perform the following steps to prevent Tru64 AdvFS from losing access to DS6000 disk drive modules, as a result of temporary path loss. For Tru64 UNIX 5.1B, you must change the **AdvfsIORetryControl** parameter. The default value is 0, and the range could be from 0 to 9.

1. Edit the sysconfig file and increase the number of AdvFS I/O retries, for example from 0 to 2.

```
# /sbin/sysconfig -r advfs AdvfsIORetryControl=2
AdvfsIORetryControl: reconfigured
```

2. Perform the following to prevent the parameter value from reverting to the default value of 0 after a system restart:

```
# sysconfig -q advfs AdvfsIORetryControl > /tmp/advfs.out
# vi /tmp/advfs.out
advfs:
AdvfsIORetryControl=2
# sysconfigdb -af /tmp/advfs.out advfs
-> New entry in the /etc/sysconfigtab
# sysconfig -d advfs
advfs:
AdvfsIORetryControl=2
```

IBM eServer iSeries host attachment

This chapter describes how you attach an IBM eServer iSeries™ host system to a DS6000 with fibre-channel adapters.

This chapter contains the following host attachment topics for an IBM eServer iSeries:

- Locating the WWPN for an IBM eServer iSeries host
- LUN considerations for IBM iSeries hosts
- Switch support for IBM iSeries hosts
- Recommended configurations for IBM iSeries hosts
- Running the Linux operating system on an IBM i5 host
 - Running Linux in a guest partition on an IBM i5 server
 - Planning to run Linux in a hosted or nonhosted guest partition
 - Creating a guest partition to run Linux
 - Managing Linux in a guest partition
 - Ordering or upgrading a server to run a guest partition

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Notes:

1. You cannot use IBM System Storage Multipath Subsystem Device Driver on the IBM iSeries host system. Multipathing runs on the i5/OS. You must be running with a minimum operating system level of i5/OS Version 5 Release 3 (V5R3) or later, to be able to take advantage of multipathing.
2. With i5/OS V5R3 or later, you can assign storage unit LUNs to multiple iSeries fibre-channel adapters through switches, direct connection, or through a fabric. These multiple adapters must all be within the same i5/OS LPAR.
3. With a pre-i5/OS V5R3, iSeries system, you cannot share an IBM iSeries volume with more than one fibre-channel system attachment.
4. You can obtain documentation for the IBM iSeries host system from <http://publib.boulder.ibm.com/series/>.

Locating the WWPN for an IBM eServer iSeries host

Complete this task to locate the WWPN for an IBM iSeries host.

Note: This section applies only to i5/OS partitions.

You can find the 16-digit worldwide port name (WWPN) for an iSeries host system in the field for port worldwide name by displaying the detail on the 2766, 2787 or 5760 IOA Logical Hardware Resource information in the Hardware Service Manager in SST/DST. You can also find the WWPN by appending the number 1000 to the beginning of the 12-digit IEEE address found on the tailstock label of the fibre-channel input/output adapter (IOA).

You can locate the WWPN for an iSeries host by performing the following steps:

1. On the iSeries Main Menu panel, type strsst.
2. On the Start Service Tools (STRSST) Sign On panel, type your service tools user ID and password.
3. On the System Service Tools (SST) panel, type 1 to select **Start a service tool**.
4. On the Start a Service Tool panel, type 7 to select **Hardware service manager**.
5. On the Hardware Service Manager panel, type 1 to select **Packaging hardware resources (systems, frames, cards,...)**.
6. On the Packaging Hardware Resources panel, type 9 to select the **System Expansion unit**.
7. On the Packaging Hardware Resources panel, type 8 to select **Storage IOA**.
8. On the Logical Resources Associated with a Packaging Resource panel, type 5 to select **Storage IOA**.
9. On the Auxiliary Storage Hardware Resource Detail panel, locate the field name for the worldwide port name. The number in the right column is the WWPN.

Note: If you have exchanged a 2766 fibre-channel IOA in the iSeries host, you must update the WWPN of the new 2766, 2787 or 5760 IOA.

LUN considerations for IBM iSeries hosts

This section describes LUN considerations for IBM iSeries hosts.

The storage unit creates logical unit numbers (LUN) that are nine characters. For example: 75-1409194, where 75 is the manufacturer, 14 is the logical subsystem, 09 is the volume, and 194 is the storage serial number (or WWNN)

Notes:

1. Fibre-channel attached LUNs are identified as the storage unit device type of 1750, on the IBM iSeries host system.
2. You can specify 1 - 32 LUNs for each attachment to an IBMiSeries fibre-channel adapter feature 2766, 2787, or 5760.

The following is an example of the output of the hardware service manager auxiliary storage hardware resource detail for the 2766 adapter. Similar information is displayed for the 2787, 5760, and adapters.

Description	: Storage IOA
Type-model	: 2787-001
Status	: Operational
Serial number	: 1F-A449F3B
Part number	: 0000080P6454
Resource name	: DC03
Worldwide port name	: 10000000C942B88C
PCI bus	:
System bus	: 35
System board	: 0
System card	: 32
Storage	:
I/O adapter	: 2
I/O bus	:
Controller	:
Device	:

Switch support for IBM iSeries hosts

This section contains references to supported switches for IBM iSeries hosts. This section applies only to i5/OS partitions.

The IBM iSeries supports only a homogeneous environment (only IBM iSeries initiators). You can establish a homogeneous environment by using the logical zoning of the switch. All host systems within an IBM iSeries zone must be IBM iSeries systems.

For a list of supported switches for an IBM iSeries host see: <http://www-1.ibm.com/servers/storage/support/san/>.

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Recommended configurations for IBM iSeries hosts

This section provides recommended configurations for IBM iSeries hosts. This section only applies to i5/OS partitions.

You can use the following configurations for each feature code:

- For feature codes 2766 and 2787 (2 Gb adapters):
 - Install feature code 2766, which is an I/O adapter, in the IBM iSeries system unit or in the high-speed link (HSL) PCI I/O towers.
 - Install only one 2766 adapter per I/O processor (IOP) because it requires a dedicated IOP. No other I/O adapters are supported under the same IOP.
 - Install only two 2766 adapters per multiadapter bridge.
- For feature code 5760 (4 Gb adapter):
 - Install feature code 5760, which is an I/O adapter, in the IBM iSeries system unit or in the high-speed link (HSL) PCI I/O towers.

- Install only one 5760 adapter per I/O processor (IOP) because it requires a dedicated IOP. No other I/O adapters are supported under the same IOP.
- Install only two 5760 adapters per a multiadapter bridge.

Note: 4 GB adapters can connect to a DS6000, but will auto negotiate and only run at 2 GB per second.

The following is an example of the output of the logical hardware resources of the hardware service manager that are associated with the IOP.

Opt Description	Type-Model	Status	Resource Name
Combined Function IOP	2843-001	Operational	CMB04
Storage IOA	2766-001	Operational	DC18
Disk Unit	1750-A82	Operational	DD143
Disk Unit	1750-A81	Operational	DD140
Disk Unit	1750-A81	Operational	DD101

The following is an example of the output for the hardware service manager auxiliary storage hardware resource detail for the storage unit.

Description.....	Disk unit
Type-Model.....	1750-A82
Status.....	Operational
Serial number.....	75-1409194
Part number.....	
Resource name.....	DD143
licensed machine code.....	FFFFFFFF
Level.....	
PCI bus.....	
System bus.....	35
System board.....	0
System card.....	32
Storage.....	
I/O adapter.....	6
I/O bus.....	0
Controller.....	1
Device.....	1

You can define the storage unit LUNs as either protected or unprotected. From a storage unit physical configuration view point, all IBM iSeries volumes are RAID-5 or RAID-10 volumes and are protected within the storage unit. When you create the IBM iSeries LUNs by using the IBM System Storage DS Storage Manager or DS CLI, you can create them as logically protected or unprotected.

Table 5 shows the disk capacity for the protected and unprotected models. Logically unprotecting a storage LUN, allows the iSeries host to perform remote load source mirroring to that device. Because the load source is mirrored on an external LUN, the storage unit can copy or transfer this load source as a disaster recovery backup. When you use the iSeries tool kit, an iSeries host in a remote location, using a copy of the original load source, can recover this load source and start running as if this recovery box were the original source host.

Table 5. Capacity and models of disk volumes for IBM iSeries

Size	Type	Protected model	Unprotected model	Release support
8.5 GB	1750	A01	A81	Version 5 Release 2 or later

Table 5. Capacity and models of disk volumes for IBM iSeries (continued)

Size	Type	Protected model	Unprotected model	Release support
17.5 GB	1750	A02	A82	Version 5 Release 2 or later
35.1 GB	1750	A05	A85	Version 5 Release 2 or later
70.5 GB	1750	A04	A84	Version 5 Release 2 or later
141.1 GB	1750	A06	A86	Version 5 Release 3 or later
282.2 GB	1750	A07	A87	Version 5 Release 3 or later

Starting with i5/OS Version 5 Release 3, IBM iSeries supports multipath attachment through fibre channels as part of the base i5/OS support. Version 5 Release 3 uses the existing host adapters (feature code 2766, 2787 or 5760). New paths are automatically detected, configured by the system, and immediately used. If a disk is initially set up as a single path and a second path is added, the resource name of the disk unit is modified from DDxxx to DMPxxx to reflect that it now has more than one path to the disk unit. No changes are required by the user on the iSeries to make use of the new path. Multipath connections can be made direct connect or through a fabric.

Note: Multipath is not supported on an external disk unit acting as a System i load source.

For additional information on fabric zoning recommendations see “DS6000 fabric zoning” on page 34.

To activate multipath on an iSeries host, use the IBM System Storage DS Storage Manager.

With i5/OS Version 5 Release 3 and later, path information is available for the disk unit hardware configuration from iSeries Navigator.

To improve the availability of fibre-channel disks (when there is more than one 2766, 2787 or 5760 I/O adapter), use the following configuration guidelines to help you achieve good DS6000 iSeries I/O performance:

- Spread iSeries LUNs across as many logical ranks as possible on different DS6000 processors:
 - Use no more than 500 I/O operations per second per rank.
 - Assign I/O critical iSeries applications to DDMs with speeds of 15000 rpm or more.
 - Create ranks of eight DDMs, if you are formatting RAID 5 ranks.
- Use the Disk Magic Analysis tool to calculate the smallest LUNs size as possible to allow more iSeries-active I/O, without effecting device response times.
- Assign no more than 500 GB of storage to each iSeries I/O adapter.
- For SAN attachments, do not connect more than four iSeries I/O adapters to a single I/O attachment port.
- Use only one fibre-channel IOP/IOA pair per multiadapter bridge. Do not use more than six IOP/IOA pairs per high-speed loop (HSL). Use 64-bit, 133 MHz PCI-X slots.

Running the Linux operating system on an IBM i5 server

This section describes running Linux on an IBM i5 server.

Disabling automatic system updates

Many Linux distributions give administrators the ability to configure their systems for automatic system updates. Red Hat provides this ability in the form of a program called `up2date`, and SUSE provides a program called YaST Online Update. You can configure these features to query for updates that are available for each host and to automatically install any new updates, which can include updates to the kernel.

If your host is attached to the DS6000 series and is running the IBM Multipath Subsystem Device Driver (SDD), consider turning off this automatic update feature because some drivers that are supplied by IBM, like SDD, are dependent on a specific kernel and cannot function in the presence of a new kernel. Similarly, host bus adapters drivers must be compiled against specific kernels to function optimally. By allowing automatic update of the kernel, you risk an unexpected impact to your host system.

Running Linux in a guest partition on an IBM i5 server

This section provides instructions for running the Linux operating system on IBM i5 server.

IBM and a variety of Linux distributors partnered to integrate the Linux operating system with the reliability of the i5 server. Linux brings a new generation of Web-based applications to the i5 server. IBM has modified the Linux PowerPC® kernel to run in a secondary logical partition and contributed the kernel to the Linux community. This section contains an overview of the types of tasks that you must perform to operate Linux on an i5 server:

- Plan to run Linux as a hosted or nonhosted guest partition
Determine what you need to do before you install Linux on the i5 server. Understand which software and hardware requirements that you need to support Linux. Find out the configuration options that are available and which options fit the needs of your company. Find out if the system that you own requires you to disable the multitasking function of your server's processor.
- Create a guest partition to run Linux
Understand how to configure a guest partition using system service tools (SST) and how to install Linux on the i5 server. You can also find information about supported I/O adapters (IOAs) and how to configure a network server description (NWSD.)
- Manage Linux in a guest partition
Find the information that you need to manage a guest partition running Linux. Use the information to understand which IPL types that you can use and how Linux partitions can communicate with other partitions on the server.
- Order a new server or upgrade an existing server
Use the LPAR Validation tool to validate that your planned partitions are valid. You can also find the contact information that you need to order a new server.

Go to the following Web sites for more information about running Linux on a i5 server:

- <http://www.ibm.com/servers/storage/disk/ds6000/>
- <http://publib.boulder.ibm.com/iseres/>

Planning to run Linux in a hosted or nonhosted guest partition

This section provides the planning information that you need to be able to run Linux in a hosted or nonhosted guest partition.

Evaluate each i5 server to determine if your hardware supports Linux. To successfully partition an i5 server to run Linux, the server requires specific hardware and software. The primary partition must run i5/OS Version 5 Release 3 and be updated with the latest programming temporary fixes. You can find the latest Linux-related i5/OS programming temporary fixes at <http://publib.boulder.ibm.com/series/>.

Linux is not supported on a primary partition.

Selected models can run Linux using the shared processor pool configuration. Other models require the use of dedicated processors for a Linux partition. Those same models also require that you disable processor multitasking for the whole system, including the primary partition.

The Linux operating system supports single processors or multiple processors. You make this choice when you create the guest partition. If you build a Linux kernel for a single processor that is loaded into a partition with multiple processors, the processor will function correctly. But, you can only use one processor. If you assign multiple processors to a partition, you must use Linux built for Symmetric Multiprocessors (SMP). You can assign a number of available processors to a guest partition.

To determine whether your system supports Linux in a guest partition, go to <http://publib.boulder.ibm.com/series/>.

Creating a guest partition to run Linux

These sections provide the information that you need to configure a guest partition and how to install Linux on an i5 server.

You can also find information about supported input/output (I/O) adapters and about network server configuration.

Hosted versus nonhosted guest partition running Linux

A hosted guest partition is dependent on a hosting i5/OS® partition for I/O resources. The hosting i5/OS partition can either be a primary or a secondary partition. The I/O resources a guest partition can use from a hosting partition include disk, CD, and tape devices.

You must start the hosted guest partition from the hosting i5/OS partition by using a network server description (NWSD). You can use the NWSD to control the guest partition. The guest partition can only be active when the hosting partition is active and out of restricted state. When i5/OS is in a restricted state, all NWSDs are automatically varied off. An i5/OS partition can host multiple guest partitions. Ensure that the hosting partition can support guest partitions.

You can IPL a hosted guest partition by varying on an NWSD object. You must not power on a hosted guest partition when using the Work with Partitions Status display. If you power on the guest partition running Linux using the Work with Partitions Status display, all of the virtual I/O devices will be unavailable.

A nonhosted guest partition is not dependent on a hosting i5/OS partition for any I/O resources. The guest partition has its own disk units, or the partition makes use of networking support to do a network boot. You can start a nonhosted guest partition even if the primary partition is not fully active. You can start a nonhosted guest partition from the Work with Partitions Status display.

Obtaining the Linux operating system for i5 servers

Linux is an open source operating system. You can obtain Linux in source format and build it for one person or a business. The open source code encourages feedback and further development by programmers. Linux developers are encouraged to design their own specialized distribution of the operating system to meet their specific needs.

All Linux distributions share a similar Linux kernel and development library. Linux distributors provide custom components that ease the installation and maintenance of Linux systems. Before you install another distributor's version of Linux, verify that the kernel has been compiled for the Power PC and the hardware for the i5 server. Your system might be misconfigured and will not run Linux in a guest partition.

You can download different versions of Linux through the Internet. However, not all the versions of Linux have been tested for use with the storage unit. See your Linux distributor to learn how to obtain the latest maintenance updates.

Managing Linux in a guest partition

The following sections provide the information that you need to manage Linux in a guest partition.

Virtual I/O in a guest partition running Linux

Virtual I/O resources are devices that are owned by the hosting i5/OS partition. The i5 Linux kernel and i5/OS supports the following virtual I/O resources: virtual console, virtual disk unit, virtual CD, virtual tape, and virtual Ethernet.

Virtual console provides console function for the guest partition through an i5/OS partition. The virtual console can be established to the hosting partition or to the primary partition. The use of the virtual console allows the installation program to communicate with the user prior to networking resources that you configure. You can use the virtual console to troubleshoot system errors.

When you use virtual disk for a Linux partition, the i5/OS partitions control connectivity to the real disk storage. The i5/OS hosting partition and its operating system versions control the storage unit connectivity solely in this configuration.

A hosting partition can only provide a virtual disk unit. Virtual DASD provides access to NWSSTG virtual disks from Linux. By default, the CRTNWSSTG command creates a disk environment with one disk partition that is formatted with the FAT16 file system. The Linux installation program reformats the disk for Linux, or you can use Linux commands such as **fdisk** and **mke2fs** to format the disk for Linux.

Notes:

1. For more information about i5/OS connectivity to the storage unit, see <http://www-1.ibm.com/servers/storage/disk/ds6000>
2. For more detailed information about how your company might use a guest partition with I/O resources go to <http://publib.boulder.ibm.com/iseres/>.

Directly attached I/O in a guest partition running Linux

With directly attached I/O, Linux manages the hardware resources and all I/O resources. You can allocate disk units, tape devices, optical devices, and LAN adapters to a guest partition running Linux.

You must have an NWSD to install Linux in a guest partition. After you install Linux, you can configure the partition to start independently.

For directly attached hardware, all failure and diagnostic messages are displayed within the guest partition.

Connectivity to the storage unit from i5 Linux is managed solely through fibre-channel adapters. For more information about the adapters and the Linux device, see http://www-1.ibm.com/servers/eserver/series/linux/fibre_channel.html.

Ordering or upgrading a server to run a guest partition

This section tells you who to contact to order a new server or upgrade an existing server to run a guest partition.

The LPAR Validation tool emulates an LPAR configuration and validates the planned partitions. In addition, the LPAR Validation Tool allows you to test the placement of i5/OS and Linux hardware within the system to ensure that the placement is valid. Refer to Logical Partitions for information on the LPAR Validation Tool (LVT).

Contact your IBM marketing representative or IBM Business Partner to enter the order. You can enter the order by using the iSeries configurator. The configurator has been enhanced to support ordering IOAs without IOPs when you define a Linux partition.

IBM eServer pSeries or IBM RS/6000 host attachment

This chapter describes how you attach an IBM eServer pSeries™ or IBM RS/6000® host system to a DS6000 with fibre-channel adapters.

This chapter contains the following sections:

- Installing the host attachment package on IBM pSeries AIX hosts
- Locating the WWPN for an IBM eServer pSeries or an RS/6000 host
- Attaching a pSeries or RS/6000 host
- Support for fibre-channel boot on a pSeries host
- Attaching to multiple RS/6000 or pSeries hosts without the HACMP host system
- Installing the Subsystem Device Driver (SDD) on an IBM pSeries host running Linux

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

- The IBM® System Storage™ Multipath Subsystem Device Driver User's Guide at <http://www.ibm.com/systems/support/storage/software/sdd>.

Notes:

1. Ensure that you can access the installation script files. These files are on the CD that you receive with the storage unit.
2. Ensure that you have 1 MB minimum of available hard disk space to install the AIX host attachment package.
3. The SDD supports RS/6000 and pSeries host systems in a clustering environment. To have failover protection on an open system, SDD requires a minimum of two paths, and SDD allows a maximum of 32 paths.
4. The storage unit supports eServer BladeCenter JS20 with blade 8842 in an IBM pSeries, 32-bit environment.
5. For an RS/6000 or pSeries host system, you can use either the switched fabric or arbitrated loop topology. The RS/6000 and pSeries host systems *do not support* more than one host bus adapter on the loop.

Installing the host attachment package on IBM pSeries AIX hosts from the CD

Complete this task to install the host attachment package on IBM pSeries AIX hosts from the CD.

Before installing the host attachment package, you must comply with the following prerequisites.

1. You must have superuser authority and administrator knowledge to complete the instructions.
2. You must have knowledge of the System Management Interface Tool (SMIT).
3. Your host system must be attached to a DS6000.

Notes:

1. The following example uses /dev/cd0 for the address of the CD drive. Your address might be different.
2. For information about upgrading the host attachment package on an IBM pSeries AIX hosts, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* at <http://www.ibm.com/systems/support/storage/software/sdd>.

Complete the following steps to install the host attachment package on an IBM pSeries AIX host:

1. Turn on the host system and all attachments.
2. From your AIX prompt, type `smit install_update` to go directly to the installation panel.
3. Highlight **Install Software** and press **Enter**.
4. Press **F4** to open the Input Device/Directory for Software window.
5. Select the CD drive that you are using for the installation; for example, /dev/cd0.
6. Press **Enter**. The Install and Update from the Latest Available Software window opens.
7. Highlight **Software to Install** and press **F4**.
8. Select **Software Packages** and press **F7**.

The software package name is `devices.fcp.disk.ibm.rte` or `devices.fcp.disk.ibm.mpio.rte`.

9. Press **Enter**. The Install and Update from the Latest Available Software panel opens again with the name of the selected software highlighted.
10. Check the default option settings to ensure that they are what you need.
11. Press **Enter** to install the software. SMIT responds with the following question: Are you sure?
12. Press **Enter** to continue. The installation process might take several minutes. A status bar will report **ok** when the process is complete.
13. Press **F10** when the installation process is complete.
14. Exit from SMIT.
15. Remove the CD.
16. You must restart the host system for activation.

Installing the host attachment package on IBM pSeries AIX hosts from the SDD Web site

Complete this task to install the host attachment package on IBM pSeries AIX hosts from the IBM System Storage Multipath Subsystem Device Driver (SDD) Web site.

Before installing the host attachment package, you must comply with the following prerequisites.

1. You must have superuser authority and administrator knowledge to complete the instructions.
2. You must have knowledge of the System Management Interface Tool (SMIT).
3. Your host system must be attached to a DS6000.

Note: For information about upgrading the host attachment package on an IBM pSeries AIX hosts, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* at <http://www.ibm.com/systems/support/storage/software/sdd>.

Complete the following steps to install the host attachment package on an IBM pSeries AIX host:

1. Turn on the host system and all host attachments.
2. Download the host attachment software at <http://www.ibm.com/systems/support/storage/software/sdd>.
3. Use FTP to copy the software to the applicable host.
4. If the download package is in .tar format, issue the `tar -xvftar` command to untar the package. See <http://www-1.ibm.com/support/docview.wss?uid=ssg1S1001580> for more information.
5. From your AIX prompt, type `smit install_update` to go directly to the installation panel.
6. Highlight **Install Software** and press **Enter**.
7. Type in the location of software. (For example, `/tmp`).
8. Press **Enter**. The **Install and Update from the Latest Available Software** window opens.
9. Highlight **Software to Install** and press **F4**.
10. Select **Software Packages** by highlighting them and pressing **F7**.
11. Check the default option settings to ensure that they are correct.

12. Press **Enter** to install the software. SMIT responds with the following question:
Are you sure?
13. Press **Enter** to continue. The installation process might take several minutes. The status bar indicates when the process is complete.
14. Press **F10** when the installation process is complete to exit SMIT.
15. You must restart the host system for the changes to take effect.

Locating the WWPN for an IBM eServer pSeries or an RS/6000 host

Complete this task to locate the WWPN for an IBM pSeries or RS/6000 host.

1. Log in as root.
2. Type `lscfg -vl fcsx`, where *x* is the adapter number.
The network address is the fibre-channel adapter port WWPN value.

Note: The `lscfg -vl fcsx` ROS level identifies the fibre-channel adapter firmware level.

To display all adapters in the systems with the associated WWPN when you are using SDD, enter the following SDD command: `datapath query wwpn`.

Attaching a pSeries or RS/6000 host

This section provides instructions, requirements, and considerations for attaching an RS/6000 or pSeries host system to a storage unit with fibre-channel adapters.

Verifying the AIX host system configuration on a pSeries

This section provides a command you can run on an AIX host system to verify the configuration of the storage unit for fibre-channel adapters on a pSeries host.

To verify the configuration of the storage unit on a host system running AIX, type the following command:

```
lsdev -Cc disk | grep 1750
```

The following is example output if the configuration is successful:

```
hdisk3 Available 30-68-01 IBM FC 1750
hdisk4 Available 30-68-01 IBM FC 1750
hdisk5 Available 30-68-01 IBM FC 1750
...
...
```

The following is example output from the `lsdev -Cc disk | grep Other` command, which lists all the disks in the systems that do not contain 1750 disk devices.

```
hdisk3 Available 30-68-01, Other FCSCSI disk device
hdisk4 Available 30-68-01, Other FCSCSI disk device
hdisk5 Available 30-68-01, Other FCSCSI disk device
...
...
```

Troubleshooting a host running AIX:

This section provides a troubleshooting configuration aid for a DS6000 that is connected to a host system that is running AIX.

On a heavily loaded system, you might see lower than expected performance and errors that indicate that the host adapter was unable to activate an I/O request on the first attempt. The most likely cause of these errors is that the host is low on resources.

To reduce the incidence of these errors, perform the following steps to increase the resources by modifying the maximum transfer size attribute for the adapter:

1. Type the following command to view the current setting:
lsattr -El <hba> -a max_xfer_size, where <hba> is the name of the adapter that logs the error.
2. Type the following command to increase the size of the setting:
chdev -l <hba> -P -a max_xfer_size=0x1000000

Note: To view the range of allowable values for the attribute, type **lsattr -Rl <hba> -a max_xfer_size**.

3. Restart the host to put these changes into effect.

Making SAN changes for IBM pSeries hosts

Before you change the switch port that is connected to the storage unit or the switch domain ID, perform the steps in this section to remove the fibre-channel adapter.

Note: The 1905, 1910, 1957, 1977, 5716, 5758, 5759, 5773, 5774, 6228, and 6239 host adapters do not support dynamic SAN configuration changes. Do not change the domain ID of the switch port of the switch that is connected to the storage unit; otherwise, you might lose the connection to the storage unit.

Complete the following steps on a pSeries host to remove the fibre-channel adapter before changing a switch port:

1. Run the **umount** command to unmount all applicable file systems.
For example, type **umount x**, where *x* is the file system name. If you are on an SP host, go to step 2; otherwise, go to step 4.
2. For SP™ environments, type **suspendvsd** to suspend the virtual shared disks with the volume groups.
3. For SP environments, type **stopvsd** to stop the virtual shared disks that are associated with the volume groups.
4. Run the **varyoffvg** command for the 1750 volume group.
For example, type **varyoffvg VGname**.
5. Type **rmdev -dl fcs(x) -R** on the command line to remove the fibre-channel adapter. The value of *x* represents the fibre-channel number that you want to remove.
6. Make all required San changes.
7. Type **cfgmgr** to reinstall the adapter and rediscover the 1750 hdisks.
8. Type **varyonvg** to vary on the volume groups.
9. Run the **mount** command to remount all applicable file systems.

Support for fibre-channel boot on a pSeries host

The following sections describe the prerequisites, considerations, installation mechanisms, and disk configuration types for a fibre-channel boot operation on a pSeries host.

You must have the correct level of firmware, microcode, and operating system for your adapter and host. For the most current information about hosts, feature codes, firmware, microcode, adapters and switches for fibre-channel boot that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/>.

Retrieve a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>.

For additional SDD information see, *IBM System Storage Multipath Subsystem Device Driver User's Guide*. You can find this guide and the latest driver download at <http://www.ibm.com/systems/support/storage/software/sdd>.

Notes:

1. You must get the microcode from an IBM Web site.
2. RS/6000 supports fibre-channel boot with or without SDD. However, if you are using SDD with RS/6000 SP, you must install APAR IY37451.
3. If you are using AIX 5.1 the SP, PSSP must be at 3.2 with APAR IY22732 (included in PSSP 3.2 PTF 14).
4. Before configuring a pSeries host for boot-support, ensure you have copy of the **devices.fcp.disk.ibm.rte** file.
5. You must define the root volume group (rootvg) for each host. When a volume group is varied online, a reserve is placed on all disks in the group. For rootvg, this group is not released. You cannot share a rootvg volume between hosts even when you shut down all the hosts except one.
6. You can create only one path to the rootvg volume during the installation.
7. Mirror the storage unit rootvg so that base functions and reliability of the operating system for the storage unit are not affected.
8. If you use the multiboot, scanning for bootable disks takes much longer with a large number of accessible storage elements in the Storage Area Network.

Fibre-channel boot considerations for IBM pSeries hosts

This section lists fibre-channel boot considerations for IBM pSeries hosts.

Before you plan to use connect a fibre-channel host, consider the following:

- You must define the rootvg for each host.
When a volume group is varied online, a reserve is placed on all disks in the group. For rootvg, this group is not released. You cannot share a rootvg volume between hosts even when you shutdown all the hosts except one.
- You can create only one path to the rootvg volume during the installation.
- If you use the multiboot, scanning for bootable disks takes much longer with a large number of accessible storage elements in the Storage Area Network.
- Mirror the root volume group (rootvg)
By mirroring the storage unit root volume group, the base function and reliability of the operating system for the storage unit is not affected.

pSeries host boot device installation and disk configuration

This section provides installation mechanisms and configuration methods for boot installation from fibre-channel SAN DASD on IBM pSeries hosts.

You can use the following mechanisms to direct a PSSP boot installation on an IBM RS/6000 SP node:

- SAN_DISKID, which is derived from the worldwide port name and LUN ID

Note: This is the preferred mechanism.

- PVID
- HDISK

You can use the following methods to configure a disk for the IBM RS/6000 for a boot install from fibre-channel SAN DASD:

- Mirrored volume
- Alternate volume
- Individual volume

Attaching to multiple RS/6000 or pSeries hosts without a HACMP host

This section provides the instructions for attaching one or two storage units to multiple host systems without the high availability cluster multi-processing (HACMP™) clustering solution.

Considerations for multiple RS/6000 or pSeries hosts without HACMP

This section provides considerations for attaching a DS6000 to multiple RS/6000 or pSeries hosts without the HACMP host system.

You can install HACMP to define and access a unique journaled file system (JFS) file that is stored on a single storage unit from any attached host system. If you choose not to install HACMP and are attaching multiple RS/6000 or pSeries hosts, consider the following:

- Without HACMP, some system failure management features such as failover are not available. Therefore, a failure on the storage unit or any one of the connected host systems will most likely affect the availability of the other connected devices.
- Multiple host systems cannot access the same volume group or the same journaled file system simultaneously.
- You must vary on and mount the volume groups and journaled file systems every time that you start the system.
- The storage unit does not allow ownership of volume groups to move from one system to another.
- When you use this procedure, you can define between 2 - 4 multiple host systems.

Saving data on the storage unit when attaching multiple RS/6000 or pSeries hosts

You can save data on the storage unit when you attach multiple RS/6000 or pSeries host systems to the storage unit.

Note: You can restore the data that you save using this procedure by following the steps in “Restoring data when attaching multiple RS/6000 or pSeries” on page 81.

Perform the following steps to preserve the data that is stored on a storage unit that was previously connected to a RS/6000 or pSeries host. This procedure does not erase the data on the storage unit, but it removes the volume groups from the host system:

1. Type `umount` to unmount all file systems from all host that are connected to the storage unit.
2. Type `fsck` on each of the file systems on the storage unit to verify the file system integrity.
3. Type `varyoffvg` to vary off all the storage unit volume groups from all of the host systems that are connected to the storage unit.
4. Type `exportvg` to remove all of the storage unit volume groups from all the host systems that are connected to the storage unit.
5. Type `rmdev -ld hdiskx` to delete all physical volumes (hdiskx) on each host system that is associated with the storage unit.

Restoring data when attaching multiple RS/6000 or pSeries

You can restore data on the storage unit when you attach multiple RS/6000 or pSeries host systems to the storage unit.

This procedure assumes that you have preserved the data by following the instructions in “Saving data on the storage unit when attaching multiple RS/6000 or pSeries hosts” on page 80.

Complete these steps to restore the data:

Note: You need to complete the steps for each host system.

1. Check to be sure that the storage unit physical volumes (hdiskx) are available.
Type `lsdev -Cc disk` to display the hdiskx on the host system.
2. Type `importvg -y xxxx zzzz` where `xxxx` is the name of the volume group, and `zzzz` is the physical volume that you want to restore.
3. Type `smit chvg` to verify that **No** is selected for the following two options:
 - **Activate Volume Group Automatically**
 - **A Quorum and of disks required to keep the volume group online?**
4. Type `varyonvg VGname` to vary-on the storage unit volume groups to the applicable host systems connected to the storage unit.
5. Type `mount /xxxx` where `xxxx` is the applicable file systems you want to mount.

Installing the subsystem device driver on an IBM pSeries host running Linux

This section describes considerations and requirements for installing the Subsystem Device Driver (SDD) on an IBM pSeries host running Linux. This section also provides a link for you to use to locate the most current driver and documentation.

You can install and configure the SDD using the instructions found in the *IBM System Storage Multipath Subsystem Device Driver User's Guide*. You can find this guide and the latest driver download at <http://www.ibm.com/systems/support/storage/software/sdd>.

Before you install the SDD on an IBM pSeries host running Linux, review the following considerations and also review “Hardware requirements for the subsystem device driver on a pSeries host running Linux” on page 82.

- Many Linux distributions give administrators the ability to configure their systems for automatic system updates. Red Hat provides this in a program called up2date, while SUSE provides it as YaST Online Update. These features periodically query for updates that are available for each host, and they can be configured to automatically install any new updates. This often includes updates to the kernel.

Consider turning off this automatic update feature if you hosts are attached to the DS6000 series and run the IBM Multipath Subsystem Device Driver (SDD). Some drivers that are supplied by IBM, like the SDD, are dependent on a specific kernel and will cease to function in the presence of a new kernel. Similarly, host bus adapter (HBA) drivers must be compiled against specific kernels to function optimally. By allowing automatic updates of the kernel, you risk unexpectedly impacting your host systems.

- If you are using SUSE SLES 8 or 9 or Red Hat Enterprise Linux 3.0 or 4.0, you must have Service Refresh 1 to ensure that the SDD installs and operates successfully.
- SLES 10 is supported with Device Mapper Multipath (DMM). For additional information see: <http://www.ibm.com/systems/support/storage/software/sdd>. Click on the download tab, then select the System Storage Multipath Subsystem Device Driver downloads link. Locate the Subsystem Device Driver for Linux link, and select it. This will take you to the SLES 10 Device Mapper Multipath Configuration File for DS8000 / DS6000 / ESS.
- SDD does not support the following environments for the Linux operating system on a pSeries host:
 - A host system with fibre-channel connection to a shared storage unit LUN
 - Logical volume manager
 - System start from an SDD pseudo device
 - EXT3 filesystem on an SDD pseudo device
 - System paging file on an SDD pseudo device
 - Single-path mode during concurrent download and installation of LMC
 - Single-path mode during any maintenance that impacts the path attachment, such as replacing a storage unit host adapter

Hardware requirements for the subsystem device driver on a pSeries host running Linux

This section provides the list of hardware that you need to ensure that the Subsystem Device Driver (SDD) installs and operates successfully on a pSeries host running Linux.

You must have the following hardware to ensure that the SDD installs and operates successfully:

- A pSeries host attached to a DS6000 with one or more pairs of fibre-channel host adapters

Notes:

1. To use the I/O load balancing and failover features of the SDD, you must have a minimum of two paths. Go to <http://www.ibm.com/servers/storage/disk/ds6000/> for more information about the fibre-channel adapters that you can use on your Linux host system.
2. Retrieve a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>.

- Subsystem LUNs that you created and configured for multiport access. Each LUN can have up to 32 SCSI disk instances, one for each path on the server.
- A fibre optic cable to connect each fibre-channel adapter to a DS6000 port or to two switch ports subsequently zoned to DS6000 ports.

Preparing to install the subsystem device driver on a pSeries host running Linux

This section provides the information you need to prepare to install the subsystem device driver (SDD) for the Linux operating system on your pSeries host.

You must configure the storage unit for your host system and attach the required fibre-channel adapters before you can install the SDD.

Configuring the storage unit

Configure your storage unit for multiport access for each LUN before you install the SDD. To use the load balancing and path failover protection features, the SDD requires a minimum of two independent paths that share the same LUN. Failover protection is not provided with a single path. A host system with a single fibre-channel adapter connected through a switch to multiple storage unit ports is considered a multipath fibre-channel connection.

See *IBM System Storage DS6000 Introduction and Planning Guide* for more information about configuring the storage unit for the SDD.

Configuring the fibre-channel adapters for the storage unit

You must configure the fibre-channel adapters and the adapter drivers that are attached to your host system before you can install the SDD. Follow the adapter configuration instructions to configure the adapters. See your Linux on pSeries documentation for details about installing the device driver on your fibre-channel adapter.

Installing the Subsystem Device Driver on a pSeries host running Linux

This section provides the steps that you need to complete to install the Subsystem Device Driver (SDD) on a pSeries host running the Linux operating system.

Before you install the SDD, make sure that your DS6000 is configured for multiport access for each LUN. To use the load balancing and path failover protection features, SDD requires a minimum of two independent paths that share the same LUN. Failover protection is not provided with a single path. A host system with a single fibre-channel adapter connected through a switch to multiple storage unit ports is considered a multipath fibre-channel connection. See *IBM System Storage DS6000 Introduction and Planning Guide* for more information about configuring the storage unit for the SDD.

Note: See “Upgrading the Subsystem Device Driver for the Linux operating system on the pSeries host” on page 84 for information on upgrading the SDD.

Complete these steps to complete the installation of the SDD:

1. Log onto your host system as the root user.
2. Insert the SDD installation CD into your CD-ROM drive.
3. Type `mount /dev/cdrom /mnt/cdrom` to mount the CD-ROM drive.

4. Type `cd /mnt/cdrom` to change to the `/mnt/cdrom` directory.
5. For SUSE LINUX, type `cd suse` to change to the `suse` directory. Then type `ls` to display the name of the package.
6. Type `rpm -qpi IBMsdd-1.4.0.0-6.ppc64.sles9.rpm` to view all the files in the package.
7. Type `rpm -iv IBMsdd-1.4.0.0-6.ppc64.sles9.rpm` to install the SDD. This displays a message similar to the following message:

```
Preparing for installation ...
IBMsdd-1.4.0.0-6
```

Upgrading the Subsystem Device Driver for the Linux operating system on the pSeries host

This section provides steps you must complete to upgrade the SDD for the Linux operating system on the pSeries host.

This section provides steps you must complete to upgrade the Subsystem Device Driver (SDD) for the Linux operating system on the pSeries host. You must have a preexisting version of the SDD installed on your host before you can perform these steps. See “Installing the Subsystem Device Driver on a pSeries host running Linux” on page 83 if you are installing the SDD for the first time.

1. Log onto your host system as the root user.
2. Insert the SDD installation CD into your CD-ROM drive.
3. Type `mount /dev/cdrom /mnt/cdrom` to mount the CD-ROM drive.
4. Type `cd /mnt/cdrom` to change to the `/mnt/cdrom` directory.
5. For SUSE LINUX, type `cd suse` to change to the `suse` directory. Then type `ls` to display the name of the package.
6. Type `rpm -qpi IBMsdd-1.4.0.0-6.ppc64.sles9.rpm` to view all the files in the package.
7. Type `rpm -U IBMsdd-1.4.0.0-6.ppc64.sles9.rpm` to upgrade the SDD. This displays a message similar to the following message:

```
Preparing for installation ...
IBMsdd-1.4.0.0-6
```

Verifying the subsystem device driver on a pSeries host running Linux

This section provides the steps to verify the SDD installation on the Linux operating system for your pSeries host.

Complete the following steps to verify the SDD installation:

1. Check the `/opt/IBMsdd` directory for the SDD driver and its major components. The SDD installs the device driver and its utilities. The following table lists the SDD driver and its major components.

Filename	Location	Description
<code>sdd-mod.o-xxx</code>	<code>/opt/IBMsdd</code>	The SDD device driver file, where <code>xxx</code> represents the kernel level of your host system and <code>y</code> represents <code>smp</code> or <code>up</code> .
<code>vpath.conf</code>	<code>/etc</code>	The SDD configuration file.

Filename	Location	Description
executables	/opt/IBMsdd/bin	The SDD configuration and status tools.
	/usr/bin	The symbolic links to the SDD utilities.
sdd.rcscript	/etc/init.d/sdd	The symbolic link for the SDD system startup option.
	/usr/sbin/sdd	The symbolic link for the SDD manual start or restart option.

2. Issue the **rpm -qi IBMsdd** command to receive information on the particular package. You can also issue the **rpm -ql IBMsdd** command to list the specific SDD files that you successfully installed on your Linux host system.
3. If the installation was successful, issue the **cd /opt/IBMsdd** command and then the **ls -l** command to list all of the installed SDD components. This creates output similar to the following screen:

```
total 580
-rw-r----- 1 root    root      8422 Sep 26 17:40 LICENSE
-rw-r----- 1 root    root      9120 Sep 26 17:40 README
drw-r----- 2 root    root      4096 Oct  2 16:21 bin
-rw-r----- 1 root    root     88190 Sep 26 17:40 sdd-mod.o-2.4.19
-rw-r----- 1 root    root     88817 Sep 26 17:40 sdd-mod.o-2.4.19-smp
-rw-r----- 1 root    root     88094 Sep 26 17:40 sdd-mod.o-2.4.19
-rw-r----- 1 root    root     88689 Sep 26 17:40 sdd-mod.o-2.4.19-smp
-rw-r----- 1 root    root     88775 Sep 26 17:40 sdd-mod.o-2.4.19
-rw-r----- 1 root    root     89370 Sep 26 17:40 sdd-mod.o-2.4.19-smp
```

4. The installation process packages the SDD utilities as executable files in the /bin directory. Issue the **cd /opt/IBMsdd/bin** command and then the **ls -l** command to view the files in the /opt/IBMsdd/bin directory. This creates output similar to the following screen:

```
total 232
-rwxr-x--- 1 root    root     32763 Sep 26 17:40 cfgvpath
-rwxr-x--- 1 root    root     28809 Sep 26 17:40 datapath
-rwxr-x--- 1 root    root       1344 Sep 26 17:40 sdd.rcscript
-rwxr-x--- 1 root    root     16667 Sep 26 17:40 lsvpcfg
-rwxr-x--- 1 root    root     78247 Sep 26 17:40 pathtest
-rwxr-x--- 1 root    root     22274 Sep 26 17:40 rmvpath
-rwxr-x--- 1 root    root     92683 Sep 26 17:40 addpaths
```

If the installation failed, the system displays a message similar to the following message:

```
package IBMsdd is not installed
```

Configuring the Subsystem Device Driver

See *IBM System Storage Multipath Subsystem Device Driver User's Guide* at <http://www.ibm.com/systems/support/storage/software/sdd> for information about configuring and using the SDD.

IBM eServer zSeries or IBM S/390 host attachment

This chapter describes how you attach a zSeries or S/390 host system to a DS6000 with a fibre-channel or FICON adapter.

This chapter contains the following sections:

- FICON-attached S/390 and zSeries hosts overview

- FICON migration overview for zSeries hosts
- Linux for S/390 and zSeries
- Registered state-change notifications (RSCNs) on zSeries hosts
- Analyzing service information messages for S/390 and zSeries systems

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

FICON-attached S/390 and zSeries hosts overview

This section describes how you can attach the DS6000 storage unit to FICON-attached S/390 and zSeries hosts.

Each storage unit fibre-channel adapter has four ports. Each port has a unique worldwide port name (WWPN). You can configure the port to operate with the FICON upper-layer protocol. For FICON, the fibre-channel port supports connections to a maximum of 128 FICON hosts. On FICON, the fibre-channel adapter can operate with fabric or point-to-point topologies. With fibre-channel adapters that are configured for FICON, the storage unit provides the following configurations:

- Either fabric or point-to-point topologies
- A maximum of 506 logins per fibre-channel port
- A maximum of 2000 logical paths on each fibre-channel port
- Access to all 64 control-unit images (16 384 CKD devices) over each FICON port
- 1750 Model 511/522
- Logical Subsystems 32
- Logical Volumes 8192
- Volume Groups 1040

Note: FICON host channels support more devices than the 4096 possible devices on a storage unit. This enables you to attach other control units or other storage units to the same host channel up to the limit that the host supports.

The storage unit supports the following operating systems for S/390 and zSeries hosts:

- Transaction Processing Facility (TPF)
- Virtual Storage Extended/Enterprise Storage Architecture (VSE/ESA™)
- zSeries
- z/VM®
- z/VSE™
- Linux

For details about models, versions of operating systems, and releases that the storage unit supports for these host systems, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/>.

Attaching a zSeries host with FICON adapters

This section describes how to configure a zSeries host for FICON attachment to a IBM System Storage DS6000.

You can perform FICON channel attachment on the IBM System Storage DS6000. See the *IBM System Storage DS6000 Introduction and Planning Guide* for a list of the FICON host adapter features codes. This publication also contains the number of FICON host adapters, cable group numbers, number of cables, and connector IDs to order for the cables.

FICON attachment overview for a zSeries or S/390 host:

This section provides an overview of FICON channel attachment on a S/390 or zSeries host.

You can use either longwave or shortwave adapters with the IBM S/390 or zSeries host systems. The adapter for the IBM S/390 has one channel port. The adapter for the zSeries has two or four channel ports.

You can attach the FICON channels directly to a storage unit or you can attach the FICON channels to a fibre-channel switch. When you attach the FICON channels directly to a storage unit, the maximum number of FICON attachments is limited by the number of installed FICON I/O ports on the storage unit. When you use a storage unit host adapter to attach to FICON channels either directly or through a switch, the I/O port is dedicated to FICON attachment. It cannot be simultaneously attached to SCSI hosts.

When you attach a storage unit to FICON channels through one or more switches, the maximum number of FICON attachments is 128 per storage unit I/O port. The directors provide very high availability with redundant components and no single points of failure or repair.

You can use the IBM 2042 Model 001 (CNT 64-port director) or IBM 2042 Model 128 (CNT 128-port director). You can use either director to attach fibre-channel hosts and devices in addition to the FICON hosts and devices. For these configurations, the fibre-channel hosts should communicate only with the fibre-channel devices. The FICON hosts should communicate only with the FICON devices. Set up zones in the directors to guarantee that none of the fibre-channel hosts or devices can affect the FICON traffic.

When you attach FICON products to switches or directors, you cannot use cascaded switches on S/390 host systems and on zSeries 800 and 900 model systems without licensed internal code (LIC) patches. You cannot configure a fabric of multiple interconnected directors and have a FICON channel that is attached to one director communicate to a FICON control unit that is attached to another director. The FICON architecture prohibits this capability. The reason for the restriction is because the base S/390 and zSeries I/O architecture uses a single byte for addressing the I/O devices. This one-byte I/O address is not compatible with the fibre-channel, 3-byte port address. The FICON solution to this problem on these host systems is to disallow switch cascading.

However, you can use cascaded switches on zSeries 800 and 900 model host systems with LIC patches installed and on all other zSeries models. To support cascaded switches, the fabric must be what is generally called a high integrity fabric. A high integrity fabric ensures data integrity by detecting any miscabling within a FICON cascaded switch fabric that could lead to I/O data being delivered

to the wrong destination. A high integrity fabric is one that both supports and is configured to use fabric-binding and insistent domain IDs. Fabric-binding is the ability of the fabric to prevent another switch from being added operationally to the fabric without the new switch being properly planned for in the fabric and for the switch to be configured with the high integrity attributes. An insistent domain ID is a function of the fabric that does not allow the switch address to be automatically changed when a duplicate switch address is added to the fabric. Instead, fabrics that use insistent domain IDs require an operator's overt action to change a switch address. The customization of fabric-binding and the setting of an insistent domain ID are normally done only at switch installation time or reinstallation time.

Note: Switches that do not support high integrity fabrics can only be used in a single-switch FICON fabric.

With fibre-channel and FICON-intermix mode, both Fibre-Channel Protocol (FCP) and FICON upper-level protocols can be supported within the same director when deployed independently by port. (Director ports operate in either fibre-channel or FICON mode). For a list of supported switches see: <http://www-1.ibm.com/servers/storage/support/san/>.

Linux for S/390 and zSeries

This section describes how to attach a storage unit to an S/390 or zSeries host system running Linux.

Before you can attach a storage unit to any host, review the requirements in the "General requirements for attaching a host system" on page 29 section.

Running Linux on an S/390 or zSeries host

This section provides the information you need to run Linux on an S/390 or zSeries host.

There are three ways to run Linux on an S/390 or zSeries host:

- Native
Linux can run on the entire machine with no other operating system.
- Logical Partition (LPAR)
You can logically partition the hardware into a maximum of 15 separate LPARs per logical channel subsystem. For example, you can have S/390 applications in one partition, VM and VSE applications in another partition, and Linux applications in another partition.
- z/VM guest support
You can run Linux as a virtual machine using z/VM. VM provides virtualization of central processing units, I/O subsystems, and memory. You can have hundreds of Linux systems running on a single S/390 or zSeries host. For example, you can offer a complete Linux server environment to each of the application developers and host production systems all on the same S/390 or zSeries host.

For SUSE SLES 9 or Red Hat Enterprise Linux 3.0, IBM supports 31- and 64-bit Linux for zSeries for a storage unit and FICON, and Fibre Channel Protocol (FCP) attached host systems.

Note: You can attach a zSeries FCP host system running SUSE SLES 9 or Red Hat Enterprise Linux 3.0 only through a switched-fabric configuration. You cannot attach the host through a direct configuration.

Attaching fibre-channel adapters to zSeries hosts running Linux

The topics in this section describe how you attach fibre-channel adapters to a zSeries host running Linux.

Attaching fibre-channel adapters on zSeries hosts running Linux:

This section provides the information that you need to attach fibre-channel adapters on zSeries hosts running Linux.

Linux solutions are available for the 31- and 64-bit environments. The availability of this option depends on the zSeries model and the Linux distribution. Fibre-channel support is available for both direct and switched attachment.

Linux for S/390 (31-bit)

Linux for S/390 is a 31-bit version of Linux. It is currently available for S/390 Parallel Enterprise Server™ G5 and G6 and the Multiprise® 3000 processors. It also runs on zSeries models in 31-bit mode. The 31-bit limitation limits the addressable main storage to 2 GB.

Linux for zSeries (64-bit)

Linux on zSeries supports the 64-bit architecture on zSeries processors. This includes all zSeries models. The 64-bit support eliminates the 31-bit storage limitation of 2 GB.

FCP support

Linux for zSeries can access FCP and FICON controllers and devices with the appropriate I/O driver support. Linux can run using the operating system that came with the system, in a logical partition or as a guest operating system under z/VM Version 4 Release 4 or later. For the most current list of supported operating system releases see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/>.

To attach storage through a fibre channel, you must have a zSeries FCP channel. An FCP channel requires the same FICON channel hardware that is used for FICON channel. You can define FICON channel hardware as either a FICON or FCP channel by specifying the correct channel path type (FC or FCP) in the IOCP channel path ID (CHPID) statement. Thus, defining FCP type channels in the IOCP reduces accordingly the total number of FICON type channels that you can configure.

Accessing the storage unit using Fibre Channel Protocol for a zSeries host:

This section describes the high level steps you must take to access the storage unit using Fibre Channel Protocol (FCP).

1. Gather the prerequisite information about the SAN setup. See “Setting up the storage unit for zSeries hosts running Linux” on page 90 for a list of the information that you must collect.

2. Set up the storage unit and configure the LUNs for use with Linux for zSeries. See “Setting up the storage unit for zSeries hosts running Linux” for the steps you must perform to set up the storage unit.
3. Manually load the required modules and parameters to access the disks. See “Viewing the LUNs for zSeries hosts running Linux” and “Setting up the Linux system to use the Fibre Channel Protocol device on zSeries hosts” on page 91 for the steps you must perform to view LUNs and set the Linux system to use FCP devices.
4. Make the disks permanently accessible. See “Adding additional permanent devices for zSeries hosts running Linux” on page 92 for the steps you must perform to add permanent devices.

Viewing the LUNs for zSeries hosts running Linux:

This section provides instructions for viewing a list of the LUNs for zSeries hosts that are running Linux.

You must collect the following software configuration information to prepare a Linux system for accessing the storage unit through a fibre channel:

- Host name of the server hosting the Linux system
- Device address and CHPID of the FCP port that is attached to the Linux machine
- Worldwide port name (WWPN) of the FCP port on the zSeries
- Fibre-channel port on the storage unit
- WWPN of the fibre-channel port on the storage unit

You can obtain that information from the hardware management console (HMC), the IBM System Storage DS Storage Manager, and the SAN switch. Figure 5 shows examples of the prerequisite information that you must obtain to run FCP Linux on the zSeries.

```
Linux hostname linuxa
CHPID of the FCP port on the z800 15
WWPN of the FCP port on the z800 50:05:07:64:01:40:01:7d
FC port on the storage unit Enclosure 3 Slot 1
WWPN of the FC port on the storage unit 50:05:07:63:00:c8:95:89
```

Figure 5. Example of prerequisite information for FCP Linux on zSeries

To view the LUNs that are the volumes assigned to your Linux system, use the IBM System Storage DS Storage Manager.

Setting up the storage unit for zSeries hosts running Linux:

This section provides instructions for setting up the storage unit for zSeries hosts running Linux.

You must first collect the following software configuration information to prepare a Linux system for accessing the storage unit through a fibre channel:

- Host name of the server hosting the Linux system
- Device address and CHPID of the FCP port that is attached to the Linux machine
- Worldwide port name (WWPN) of the FCP port on the zSeries
- Fibre-channel port on the storage unit

- WWPN of the fibre-channel port on the storage unit

You can obtain that information from the hardware management console (HMC), the IBM System Storage DS Storage Manager, and the SAN switch. Figure 6 shows examples of the prerequisite information that you must obtain to run FCP Linux on the zSeries.

```
Linux hostname linuxa
CHPID of the FCP port on the z800 15
WWPN of the FCP port on the z800 50:05:07:64:01:40:01:7d
FC port on the storage unit Enclosure 3 Slot 1
WWPN of the FC port on the storage unit 50:05:07:63:00:c8:95:89
```

Figure 6. Example of prerequisite information for FCP Linux on zSeries

To set up the LUNs for use with your Linux system, you first must define the system to the storage unit. You can do this through the IBM System Storage DS Storage Manager.

Note: These instructions assume that you are familiar with the IBM System Storage DS Storage Manager.

1. Log on to the IBM System Storage DS Storage Manager and then in the navigation select, in order, **Real-time manager** → **Manage hardware** → **Host systems**.
2. From the **Select Action** drop-down list, select **Create**, then click **Go**.
3. From the **General host information** panel, complete the following fields for each fibre-channel host adapter. When you are finished, click **OK**.
 - Host type
 - Nickname
 - Description
4. From the **Define host ports** panel, specify the host ports for this host. Click **Add** to add each host port to the defined host ports table.
5. From the **Define host WWPN** panel, specify the worldwide port names for the selected hosts. When finished, click **Next**.
6. From the **Select storage images** panel, specify the storage image for attachment. Click **Add** to add the storage image to the selected storage images table. When finished, click **Next**.
7. From the **Specify storage image parameters** panel, specify the following parameters for each storage image.
 - Host attachment identifier
 - Volume Group for host attachment
 - Type of I/O port (any valid storage image I/O port or a specific I/O port)
8. Click **Apply assignment** button to apply the current attachment assignment. Use this button to go through this page for each host attachment identifier that you want to assign to the storage image. When you are finished, click **OK**.
9. From the **Verification** panel, verify the data that represents the new host attachments. When finished, click **Finish**.

Setting up the Linux system to use the Fibre Channel Protocol device on zSeries hosts:

This section provides instructions for setting up the Linux system to use the Fibre Channel Protocol (FCP) device on zSeries hosts.

You must collect the following software configuration information to prepare a Linux system to access the storage unit through a fibre channel:

- Host name of the server hosting the Linux system
- Device address (and CHPID) of the FCP port that is attached to the Linux machine
- Worldwide port name (WWPN) of the FCP port on the zSeries
- Fibre-channel port on the storage unit
- WWPN of the fibre-channel port on the storage unit

You can obtain that information from the hardware management console (HMC), the IBM System Storage DS Storage Manager, and the SAN switch. Figure 7 shows examples of the prerequisite information that you must obtain to run FCP Linux on the zSeries.

```
Linux hostname linuxa
CHPID of the FCP port on the z800 15
WWPN of the FCP port on the z800 50:05:07:64:01:40:01:7d
FC port on the storage unit Enclosure 3 Slot 1
WWPN of the FC port on the storage unit 50:05:07:63:00:c8:95:89
```

Figure 7. Example of prerequisite information for FCP Linux on zSeries

You must have installed at least the following modules to use the FCP device with the storage unit:

- qdio — The same module is used for other qdio devices
 - scsi_mod — SCSI core
 - zfc — FCP support for zSeries Linux
 - sd_mod — SCSI disk support
1. Load the modules in the following order. You can load all modules with the **modprobe** command. You can load all modules, except **zfc**, without parameters. The **zfc** module requires parameters to map the FCP devices on the storage unit. Each device requires the following parameters:
 - a. The device number of the device that is defined in the IOCP for the FCP channel on the zSeries
 - b. The SCSI ID starting at 1
 - c. The WWPN of the storage unit fibre-channel port
 - d. The SCSI LUN within Linux starting at 0
 - e. The FCP LUN of the target volume on the storage unit
 2. Use the following format to supply the mapping parameters to the module:

```
map="devno SCSI_ID:WWPN SCSI_LUN:FCP_LUN"
map="0x0600 1:0x5005076300c89589 0:0x5301000000000000"
```

Adding additional permanent devices for zSeries hosts running Linux:

This section provides instructions for adding additional permanent devices for zSeries hosts running Linux.

You must collect the following software configuration information to prepare a Linux system to access the storage unit through a fibre-channel connection:

- Host name of the server that is hosting the Linux system
- Device address and CHPID of the FCP port that is attached to the Linux machine

- Worldwide port name (WWPN) of the FCP port on the zSeries
- Fibre-channel port on the storage unit
- WWPN of the fibre-channel port on the storage unit

You can obtain the required information from the hardware management console (HMC), the IBM System Storage DS Storage Manager, and the SAN switch.

Figure 8 shows examples of the prerequisite information that you must obtain to run FCP Linux on the zSeries.

```
Linux hostname linuxa
CHPID of the FCP port on the z800 15
WWPN of the FCP port on the z800 50:05:07:64:01:40:01:7d
FC port on the storage unit Enclosure 3 Slot 1
WWPN of the FC port on the storage unit 50:05:07:63:00:c8:95:89
```

Figure 8. Example of prerequisite information for FCP Linux on zSeries

1. Add devices using one of the following methods:

- Write a script that includes all of the parameters that are included. The following is an example script:

```
rmmod zfc
modprobe qdio
modprobe scsi_mod
insmod zfc map="\
0x0600 1:0x5005076300c89589 0:0x5301000000000000;\
0x0600 1:0x5005076300c89589 1:0x5302000000000000;\
0x0600 1:0x5005076300c89589 2:0x5300000000000000"
modprobe sd_mod
```

- Add the devices to an existing configuration with the **add_map** command. After using this command, you must manually make the devices known to the SCSI stack. The following example shows how to add SCSI devices with the **add_map** command:

```
linuxa:~#echo "0x0600 0x00000001:0x5005076300c89589 \
0x00000001:0x5302000000000000">/proc/scsi/zfc/add_map

linuxa:~#echo "scsi add-single-device 0 0 1 1">/proc/scsi/scsi
```

2. To make the devices available permanently after a reboot, you must create a new initial ramdisk (RAM that is accessed like a disk) that contains the necessary modules and parameter information. First, save the module parameters in the configuration file `/etc/zfc.conf`, as shown in the following example.

```
linuxa:~#cat /proc/scsi/zfc/map >/etc/zfc.conf
linuxa:~#cat /etc/zfc.conf
0x0600 0x00000001:0x5005076300c89589 0x00000000:0x5301000000000000
0x0600 0x00000001:0x5005076300c89589 0x00000001:0x5302000000000000
0x0600 0x00000001:0x5005076300c89589 0x00000002:0x5300000000000000
```

3. Create a new ramdisk with the **mk_initrd** utility and then run the **zipl** utility to update the initial program load record to point to the new ramdisk.

4. Enable multipath support, which is provided by the logical volume manager kernel for Linux z/FCP.
 - a. Create physical volumes with the following commands:


```
pvccreate /dev/sda1
pvccreate /dev/sdb1...
```
 - b. Put the physical volumes into a volume group with the following command:


```
vgcreate fcpvg /dev/sda1 /dev/sdb1...
```
 - c. Create as many logical volumes as you want with the following command:


```
lvcreate --size 16G fcpvg
```
 - d. Enable the alternate paths to the physical volumes with the **pvpath** command.


```
pvpath --path0 --enable y /dev/sda1
pvpath --path1 --enable y /dev/sda1
```

If both paths have a weight of 0, they will loadbalance.

As shown in the following example, these configurations yield the SCSI device /dev/sda - /dev/sdc accessed on the first path, and the SCSI device /dev/sdd - /dev/sdf accessed on the second path.

```
rmmod zfcpc modprobe qdio modprobe scsi_mod insmod zfcpc map="\ 0x0600 1:0x5005076300c89589
0:0x5301000000000000;\ 0x0600 1:0x5005076300c89589 1:0x5302000000000000;\
0x0600 1:0x5005076300c89589 2:0x5300000000000000;\ 0x0700 1:0x5005076300ca9589
0:0x5301000000000000;\ 0x0700 1:0x5005076300ca9589 1:0x5302000000000000;\
0x0700 1:0x5005076300ca9589 2:0x5300000000000000" modprobe sd_mod
```

Registered state-change notifications (RSCNs) on zSeries hosts

This section provides information about registered state-change notifications (RSCNs) on zSeries hosts.

McDATA and CNT switches ship without any configured zoning. This unzoned configuration enables the default zone on some McDATA switches. This configuration enables all ports in the switch with a fibre-channel connection to communicate with each other and to receive registered state change notifications about each other. You can set the zones. Here are some recommendations:

- If you have FICON-only environments, do not disable the default zone.
- If you have FICON environments and plan to add (or already have) fibre-channel intermix on the switch, ensure that all FICON host adapters and all FICON-capable device ports are in an explicitly defined common zone.
- If you have open-systems traffic with Fibre Channel Protocol, continue to create zones that include the host ports and all desired storage ports. Multiple overlapping zones should be used in this case to support communication from multiple host ports to share device ports.

Analyzing service information messages for S/390 and zSeries systems

Service information messages (SIMs) are generated by a storage unit for S/390 and zSeries hosts. Before installation, you can use the customization work sheets to record the following service information: the SIM severity level and the maximum

number of times the storage unit sends the SIMs (0-5 times) to the console. During installation, either you or the IBM service sales representative must enter this information into the system.

SIM message types

The following SIM message types are generated by the storage unit.

Direct access storage device (DASD) SIM

Tracks DDM failures and problems.

Media SIM

Tracks data check problems on the media.

Storage unit SIM

Tracks storage unit, power control, and other hardware problems.

SIM severity levels

1 acute

An irrecoverable error with possible loss of data. This severity level only applies to DASD SIMs.

2 serious

An irrecoverable error or a data check with loss of access to data.

3 moderate

A system path is not operational and performance might be degraded. This severity level does not apply to media SIMs.

4 service

A recoverable error, equipment checks, or data checks. You can defer repair.

IBM NAS Gateway 500 host attachment

This chapter describes how you attach an IBM NAS Gateway 500 host system to a DS6000 with fibre-channel adapters.

This chapter contains the following host attachment topics:

- Locating the WWPN on a NAS Gateway 500 host system

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Note: Feature code 6239 is a 1-port 2 gigabit fibre-channel host adapter card.
Feature code 6240 is a 2-port 2 gigabit fibre-channel host adapter card.

Locating the WWPN on a NAS Gateway 500 host system

You can obtain the worldwide port names (WWPN) of the fibre-channel adapters installed on NAS Gateway 500 through a Web browser or through the DS CLI. The following sections provide instructions for using these methods.

Obtaining NAS Gateway 500 WWPNs using a Web browser

This section provides instructions for obtaining NAS Gateway 500 WWPNs using a web browser.

If your external storage requires you to enter a worldwide port name (WWPN) for the fibre-channel host adapters that are installed in your NAS Gateway 500, you can obtain these WWPNs by using the Internet.

1. Open a Web browser.
2. Enter the following Web address: `http://hostname/NAS500GetWWN.html`
where *hostname* is the host name or IP address of your NAS Gateway 500 system. If your NAS Gateway 500 is not within the same IP subnet, use the fully qualified domain name that you use with the DNS name resolution. For example: `nasgateway500.servers.mycompany.com`.

Obtaining WWPNs through the command-line interface

Complete this task to obtain WWPNs through the command line interface.

To obtain WWPNs, perform the following steps:

1. Log in to the NAS Gateway 500 using the root user ID from a serial terminal.
2. Run the following command to install the WWPNs of all fibre-channel adapters:

```
lscfg -vpl "fcs*" |grep Network
```

The following is example output:

```
(/)-->lscfg -vpl "fcs*" |grep Network
Network Address.....10000000C93487CA
Network Address.....10000000C934863F
Network Address.....10000000C93487B8
Network Address.....10000000C934864F
```

3. You can optionally use the following command to put all vital product data of the fibre-channel adapters installed on NAS Gateway 500 into a text file:

```
lscfg -vpl "fcs*" > foo.txt
```

You can keep this file for future use. You can find information for your WWPN, location information, microcode level, part number, and other information about your fibre-channel adapters in this file.

The following is a text file example:

```

fcs2 U0.1-P2-I6/Q1 FC Adapter
Part Number.....00P4295
EC Level.....A
Serial Number.....1E323088E2
Manufacturer.....001E
Feature Code/Marketing ID...5704
FRU Number..... 00P4297
Device Specific.(ZM).....3
Network Address.....10000000C93487CA
ROS Level and ID.....02E01035
Device Specific.(Z0).....2003806D
Device Specific.(Z1).....00000000
Device Specific.(Z2).....00000000
Device Specific.(Z3).....03000909
Device Specific.(Z4).....FF601032
Device Specific.(Z5).....02E01035
Device Specific.(Z6).....06631035
Device Specific.(Z7).....07631035
Device Specific.(Z8).....20000000C93487CA
Device Specific.(Z9).....HS1.00X5
Device Specific.(ZA).....H1D1.00X5
Device Specific.(ZB).....H2D1.00X5
Device Specific.(YL).....U0.1-P2-I6/Q1
fcs3 U0.1-P2-I5/Q1 FC Adapter
Part Number.....00P4295
EC Level.....A
Serial Number.....1E3230890F
Manufacturer.....001E
Feature Code/Marketing ID...5704
FRU Number..... 00P4297
Device Specific.(ZM).....3
Network Address.....10000000C934863F
ROS Level and ID.....02E01035
Device Specific.(Z0).....2003806D
Device Specific.(Z1).....00000000
Device Specific.(Z2).....00000000

```

IBM SAN File System host attachment

This chapter describes the requirements that you need to verify before you to attach an IBM SAN File System metadata server to a DS6000 with fibre-channel adapters.

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- Support for the IBM SAN File System, at <http://www-03.ibm.com/servers/storage/software/virtualization/sfs/index.html>
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Notes:

1. You can attach a storage unit to the IBM SAN File System metadata engine using the SAN File System fibre-channel adapter. You cannot attach a storage unit to an IBM SAN File System host system using SCSI adapters.

2. You must create a host entry for each port on every SAN File System metadata engine in the cluster. For example, a 4-node SAN File System cluster requires 8 host entries. For step-by-step instructions on defining hosts and ports, see the IBM System Storage DS6000 Information Center at <http://publib.boulder.ibm.com/infocenter/ds6000ic/index.jsp> , or refer to the DS Storage Manager online help.

IBM SAN Volume Controller host attachment

This chapter provides the requirements that you need to verify before you attach an IBM SAN Volume Controller host system to a DS6000 with fibre-channel adapters.

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29,
- The *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- Support for the SAN Volume Controller host system, at , <http://www.ibm.com/servers/storage/support/software/sanvc/>
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Note: You must create a host entry for every port on every SAN Volume Controller node in your cluster. For example, a four-node SAN Volume Controller requires 16 host entries.

Intel or AMD running Linux host attachment

This chapter describes how you attach an Intel™ or Advanced Micro Devices (AMD) host running the Linux operating system to a DS6000 with fibre-channel adapters.

This chapter contains the following sections:

- Disabling automatic updates
- Downloading and installing an Emulex or QLogic adapter
- Locating the WWPN for a Linux host
- Configuring the storage unit for an Intel host running Linux
- Manually adding and removing SCSI disks
- Support for fibre-channel boot

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>
- The *IBM System Storage Multipath Subsystem Device Driver User's Guide* at <http://www.ibm.com/systems/support/storage/software/sdd>.

Notes:

1. The DS6000 supports both direct and fabric attachment to the eServer BladeCenter™ JS20 when it operates with the following releases:
 - Red Hat Enterprise Linux 3.0 or 4.0 (all releases)
 - SUSE Linux Enterprise Server (SLES) 9 SP3

See the DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for more information about BladeCenter JS20 support.

2. The DS6000 supports both single path and multipath connectivity. Multipath connectivity is only supported with the use of the IBM Multipath Subsystem Device Driver.
3. The DS6000 supports Device Mapper Multipath (DMM) for the following releases:
 - Red Hat Enterprise Linux 4.4
 - SUSE Linux Enterprise Server 9 SP3

For additional information see <http://www.ibm.com/systems/support/storage/software/sdd>. Click the download tab, and select the System Storage Multipath Subsystem Device Driver downloads link. Locate the Subsystem Device Driver for Linux link, and select it. This will take you to the SLES 10 Device Mapper Multipath Configuration File for DS8000 / DS6000 / ESS.

4. The steps in this section that are used to install and configure adapter cards are only examples. Your configuration might be different.
5. Many Linux distributions give administrators the ability to configure their systems for automatic system updates. Administrators can configure these features to query for updates that are available for each host and to automatically install any new updates that they find, which can include updates to the kernel.

If your host is attached to the DS6000 series and is running the IBM Multipath Subsystem Device Driver (SDD), consider turning off this automatic update feature because some drivers that are supplied by IBM, like SDD, are dependent on a specific kernel and cannot function in the presence of a new kernel. Similarly, host bus adapter drivers must be compiled against specific kernels to function optimally. By allowing automatic update of the kernel, you risk an unexpected impact to your host system.

Installing an Emulex adapter on an Intel or AMD host running Linux

This section describes how to install an Emulex fibre-channel adapter on an Intel or AMD host that is running the Linux operating system.

Not all Emulex adapters and drivers are supported on an Intel or AMD host that is running Linux. For details about models, versions of operating systems, and releases that the storage unit supports for these host systems, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/>.

Note: The Emulex LP9402DC adapter card is a 3.3 volt card only and requires a 3.3 volt PCI slot.

Single- and dual-port fibre-channel interfaces with an Emulex adapter card support the following public and private loop modes:

- Target
- Public initiator
- Private initiator
- Target and public initiator
- Target and private initiator

Single- and dual-port fibre-channel interfaces with an Emulex adapter card support the following fabric attach modes:

- Public loop
- Private loop
- Fabric

Note: When you are connecting through a fibre-channel switch, configure the Emulex adapter ports and fibre-channel switch ports to autonegotiate speed and topology. This allows the correct speed and topology to be configured automatically.

Perform the following steps to download and configure an Emulex adapter:

1. Record the Institute of Electrical and Electronics Engineers (IEEE) number that is printed on the card. You can use the IEEE number to determine the WWPN.

Note: If the adapter is already installed in the host system and the host system has installed a Linux distribution, it is possible to retrieve the WWPN using host system commands.

2. Verify that the Emulex adapter driver module is loaded using the host system command **lsmod**. Look for the module name beginning with **lpfc**.
3. Determine the kernel version that you are running. You can obtain the version by running the host system command **uname -r**. The version can be either 2.4x or 2.6x.
 - For 2.4 kernels, you can find the WWPN of your adapter cards by issuing the **cat /proc/scsi/adapter_type/n** command, where *adapter_type* is the host adapter type and *n* is the host adapter number for your card. The host adapter type for Emulex adapters is **lpfc**.
 - For 2.6 kernels, because of the migration to sysfs, the adapter WWPN information may not be available in **/proc/scsi/lpfc/n**. If this is the case, then in the **/proc/scsi/lpfc/** directory, record the values represented. Then **cat /sys/class/scsi_host/hostn/port_name**. Where *n* is each of the values recorded from the **/proc/scsi/lpfc/** directory.
4. Install the adapter card. If the card is not connected directly to the storage unit, or to a fabric switch port, install a loop-back connector, that is usually supplied with the adapter card, then restart the server.

Note: If the server is designed for hot plug PCI, see the server documentation for specific directions for hot plug of PCI devices.

5. Install the adapter driver. Most Linux distributions now provide Emulex adapter driver support as part of the kernel package. See the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for supported driver levels for your Emulex adapter card. If the supported level differs from the preinstalled version of the driver, then go to “Downloading and installing a host adapter driver” on page 30 or <http://www.emulex.com/ts/docoem/ibm/index.htm>. For specific adapter driver installation instructions, see the installation instructions that are provided by your host adapter vendor.

Installing and updating an Emulex driver on an Intel or AMD host running Linux

This section describes how to install and update an Emulex fibre-channel adapter driver on an Intel host running the Linux operating system.

Perform the following steps to install an Emulex adapter driver:

1. Install the adapter driver. For specific adapter driver installation instructions, see the installation instructions provided by your host adapter vendor

Notes:

- a. The *Interoperability Matrix* (at <http://www.ibm.com/servers/storage/disk/ds6000/>) provides supported driver levels for your Emulex adapter.
 - b. If the supported level differs from the preinstalled version of the driver, go to “Downloading and installing a host adapter driver” on page 30 or <http://www.emulex.com>.
 - c. Most Linux distributions now provide Emulex adapter driver support as part of the kernel package.
2. Verify that the Emulex adapter driver module is loaded using the host system command **lsmod**. Look for the module name beginning with **lpfc**.
 3. Determine the kernel version that you are running. You can obtain this version by running the host system command **uname -r**. The version can be either 2.4x or 2.6x.

Notes:

- a. For 2.4 kernels, you can find the driver version in the top line of the **/proc/scsi/lpfc/n**, where *n* is the host adapter port that is represented as a SCSI host bus in Linux. For example, *Emulex LightPulse FC SCSI 7.1.1.14*.
- b. For 2.6 kernels, because of the migration to sysfs, the driver version might not be available in the **/proc/scsi/lpfc/n** directory. If this is the case, go to the **/proc/scsi/lpfc** directory and record the values. Type **cat /sys/class/scsi_host/host n /lpfc_drvr_version**, where *n* is each of the values recorded from the **/proc/scsi/lpfc** directory.

To determine which firmware version is currently installed on your host system, perform the following steps:

Note: The *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> provides the supported driver levels for your Emulex adapter. If the supported firmware level is not correct, contact IBM hardware support to assist you in upgrading the Emulex adapter firmware.

1. Verify that the Emulex adapter driver module is loading using the host system command **lsmod**. Look for the module name beginning with **lpfc**.
2. Determine the kernel version that you are running. Obtain the version by running the host system command **uname -r**. The version can be either 2.4x or 2.6x.

Notes:

- a. For 2.4 kernels, you can find the firmware version by typing **cat /proc/scsi/lpfc/n**, where *n* is the host adapter port that is represented as a SCSI host bus in Linux. For example, *Firmware Version: 1.81 (H2D1.81X1)*

- b. For 2.6 kernels, because of the migration to sysfs, the driver version might not be available in `/proc/scsi/lpfc/n`. If this is the case, record the values that are represented in the `/proc/scsi/lpfc` directory. Then type `cat /sys/class/scsi_host/hostn/fwrev`, where *n* is each of the values recorded from the `/proc/scsi/lpfc` directory.
3. Perform one of the following steps to bring the new driver online:
 - Reboot the server. This is the most efficient option.
 - Remove and add the module using the **modprobe** Linux command.

Configuring and troubleshooting an Emulex adapter on Linux

This section describes Emulex adapter settings that can be changed based on the kernel, distribution, and IBM server hardware you are attaching to the DS6000. It also describes modifications to configuration parameters.

Before you proceed, determine the kernel version of the Linux distribution that you are running.

Type the **uname -r** command to display the kernel version that you are running.

The output displays 2.4.x or 2.6.x, where x represents the maintenance and fix level for the installed distribution.

Note: If your Emulex driver loads at system start, run the **mkinitrd** command to build a new ramdisk device to reflect this change.

For Emulex adapter cards that are running in RHEL3U5, LUN 0 might be missing from the list of available LUNS. Add the following options to the `/etc/modules.conf` file as a workaround for this issue.

1. Edit the `/etc/modules.conf` file.
2. If the value **lpfc, lpfc_inq_pqb_filter=0x1** is not present, add it and save the configuration file.
3. Use your normal procedure to reload the driver settings. For example, use the **modprobe -r** command to remove a driver and the **modprobe** to install a driver.

For Emulex adapter cards that are running in RHEL4U4 with device mapper multipath (DM-MPIO), you must add the following line to the `/etc/modprobe.conf` file if it is not already there.

1. Edit the `/etc/modprobe.conf` file.
2. If the value **options lpfc lpfc_lun_queue_depth=16** is not present, add it and save the configuration file.
3. Use your normal procedure to reload the driver settings. For example, use the **modprobe -r** command to remove a driver and the **modprobe** to install a driver.

Note: Before you unload the HBA driver, you must unmount any file systems that are using the multipath devices and any LVM volume groups that have been made unavailable, and then unload the multipath driver. Use the **umount** command to unmount the file systems. Use the **vgchange -an** command to make the LVM volume groups unavailable, the **sdd stop** command to stop the SDD multipath driver, and the **multipath -F** command to stop the DM-MPIO driver.

After you finish changing the HBA drivers and the HBA modules are reloaded, you can then restart the multipath driver, make the LVM volume groups available again, and remount the file systems. Use the **sdd start** command to start the SDD multipath driver, the **multipath** command to start the DM-MPIO multipath driver, the **vgchange -ay** command to make the LVM volume groups available again, and the **mount** command to remount the file systems.

Perform these steps to change the timeout settings for the SCSI disks to the recommended timeout value of 60:

1. Issue the following command to each SCSI LUN on the fabric: **echo "60" > /sys/block/sdXXX/device/timeout**, where **sdXXX** is each of the SCSI disks.
2. Verify the timeout value using the following command: **cat /sys/block/sdXXX/device/timeout**, where **sdXXX** is each of the SCSI disks.

Note: This setting is not persistent. You must add a script to the startup scripts to set this value correctly.

Installing a QLogic adapter on an Intel or AMD host system running Linux

This section describes how you install a QLogic fibre-channel adapter card on an Intel or AMD host system running the Linux operating system.

Notes:

1. Not all QLogic adapters and drivers are supported on all Intel hosts running Linux. For details about models, versions of operating systems, and releases that the storage unit supports for these host systems, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/>.
2. The following steps are an example configuration for a QLogic QLA23xx, and QLA246x adapter cards. The configuration for your adapter might differ.

Single- and dual-port fibre-channel interfaces with the QLogic adapter cards support the following public and private loop modes:

- Target
- Public initiator
- Private initiator
- Target and public initiator
- Target and private initiator

Single- and dual-port fibre-channel interfaces with the QLogic adapter cards support the following fabric attach modes:

- Public loop
- Private loop
- Fabric

Note: When you are connecting through a fibre-channel switch, configure the QLogic adapter ports and fibre-channel switch ports to autonegotiate speed and topology. This allows the correct speed and topology to be configured automatically.

Perform the following steps to download and configure a QLogic adapter:

1. Install the QLogic adapter card in the host system. If the card is not being connected directly to the storage unit or to a fabric switch port, install a loop-back connector, which is usually supplied with the adapter card.
2. Connect the cable to the storage unit port.
3. Restart the server.

Note: If the server is designed for hot plug PCI, see the server documentation for specific directions for hot plug PCI devices.

4. During the restart operation of the server, press **Ctrl+Q** to get to the **FAST!Util** menu panel. Use the QLogic FAST!Util utility to configure specific parameters. For instructions about using the menu system, see the QLogic documentation that is supplied with the adapter card.
5. Use the **FAST!Util** utility to set the parameters to the values that are listed in Table 6.

Table 6. Recommended settings for the QLogic adapter card for an Intel or AMD host running Linux

Parameters	Recommended settings
Link down timeout	30
Basic Input/Output System (BIOS)	Disabled (if not configuring for SAN/Remote Boot) Enabled (if configuring for SAN/Remote Boot)
Frame size	2048
Loop reset delay	5 (minimum)
Adapter hard loop ID	Disabled
Hard loop ID	0
Execution throttle	100
Fast command posting	Enabled
>4 GB addressing	Disabled (for 32-bit systems)
LUNs per target	0 or 128
Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes
Login retry count	<ul style="list-style-type: none"> • 60, if you are not running a multipath driver • 8, if you are running a multipath driver
Port down retry count	<ul style="list-style-type: none"> • 60, if you are not running a multipath driver • 8, if you are using a multipath driver
Driver load RISC code	Enabled
Enable database updates	No
Disable database load	No
IOCB allocation	256
Extended error logging	Disabled (You might want to enable this parameter if you use debugging.)

Table 6. Recommended settings for the QLogic adapter card for an Intel or AMD host running Linux (continued)

Parameters	Recommended settings
Type of connection	<ul style="list-style-type: none"> • Loop only • Point-to-point • Loop preferred (If you cannot use arbitrated loop, then default to point-to-point.) • Point-to point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop.) <p>Note: If you connect the storage unit directly to the host system, the option that you select must match the port connections on the storage unit.</p>

Installing and updating a QLogic driver on an Intel or AMD host running Linux

This section describes how to install and update a QLogic fibre-channel adapter driver on an Intel host running the Linux operating system.

Perform the following steps to install a QLogic adapter driver:

1. Install the adapter driver. See the specific adapter driver installation instructions that are provided by your host adapter vendor.

Note: The *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> provides the supported driver levels for your QLogic adapter. If the supported level differs from the preinstalled version of the driver, go to “Downloading and installing a host adapter driver” on page 30 or http://support.qlogic.com/support/oem_ibm.asp. Most Linux distributions now provide QLogic adapter driver support as part of the kernel package.

2. Verify that the QLogic adapter driver module is loaded using the host system command *lsmod*. Look for the module name beginning with *qla*.
3. Determine the kernel version that you are running. Obtain the kernel version by running the host system command *uname -r*. The version can be either 2.4x or 2.6x.

Type *cat /proc/scsi/qlxxx/n*, where *xxx* is the extension for the driver, and *n* is the host adapter port that is represented in Linux as a SCSI host bus. For example, the driver extension might be 2300, or 2xxx, and the host adapter port might be *Firmware version 3.03.02 IPX, Driver version 8.00.00b21-k*. The firmware is included with the driver version

4. Perform one of the following steps to bring the new driver online:
 - Restart the server. This is the most efficient option.
 - Remove and add the module using the *modprobe* Linux command.

Configuring and troubleshooting a QLogic adapter on Linux

This section describes QLogic adapter settings that can be changed based on the kernel, distribution, and IBM server hardware that you are attaching to the DS6000. It also describes modifications to configuration parameters that you can change if you receive errors when you are running the Linux operating system on an Intel host that is attached with a QLogic adapter.

Before you proceed, determine the kernel version of the Linux distribution that you running.

Type the **uname -r** command to display the kernel version that you are running.

The output displays 2.4.x or 2.6.x, where x represents the maintenance and fix level for the installed distribution.

For multipath solutions that use an IBM SDD or DM-MPIO, you must disable the QLogic failover mechanism. Complete these steps to disable the failover mechanism.

1. If the driver version has a **<version> -fo** extension, the failover driver is currently installed and you must disable the driver. If the driver is not installed, you do not need to complete these steps.
2. For 2.4 kernels, edit the **/etc/modules.conf** file. Add the following line to the file:
options qla2300 qla2xfailover=0.
For 2.6 kernels, edit the **/etc/modprobe.conf** or the **/etc/modprobe.conf.local** file. Add the following line to the file: **options qla2xxx qla2xfailover=0.**

For QLogic adapter cards, you must add the following line to the **/etc/modprobe.conf** (for 2.6 kernels) or to the **/etc/modules.conf** (for 2.4 kernels) file if it is not already there.

1. Edit the **/etc/modprobe.conf** file.
2. If the value **options qla2xxx ql2xmaxqdepth=16** is not present, add it and save the configuration file.
3. Use your normal procedure to reload the driver settings. For example, use the **modprobe -r** command to remove a driver and the **modprobe** command to install a driver.

For host systems that will be performing remote mirror and copy or Flash Copy functions, modify the timeout values to prevent I/O errors from occurring. The **qfull_retry_count** option (available for 2.4 kernels only) defines the number of times that a QLogic adapter card retries a command if the queue is full. The default for this option is 16.

Perform the following steps to increase the timeout value to 32.

1. For 2.4 kernels, edit **/etc/modules.conf**.
2. If the value of the **qfull_retry_count** option is not 32, change it to 32 and save the configuration file. If the option does not exist, add the following line:
options qla2xxx qfull_retry_count=32.
3. Use your normal procedure to reload the driver settings. For example, use the host system command **modprobe -r** to remove the driver and use the **modprobe** command to install the driver.

Note: If your QLogic driver loads at system start, run the **mkinitrd** command to build a new ramdisk device to reflect this change.

Locating the WWPN for a Linux host

Complete this task to locate the worldwide port name (WWPN) for an Intel server running either the Red Hat or SUSE version of the Linux operating system with a QLogic adapter.

1. To locate the WWPN without restarting the host system:
 - a. You can find the WWPN in your **/proc/scsi/qlaadapter_type/n** directory, where **adapter_type** is the host adapter type and **n** is the host adapter number for your card. Your host adapter type can be either **qlaxxxx** for QLogic adapters or **lpfc** for Emulex adapters.

2. To locate the WWPN with restarting the host system
 - a. Restart the server.
 - b. Press Alt+Q to get the **FAST!Util** menu.

Note: If you have more than one fibre-channel adapter installed, all the fibre-channel adapters display. Scroll down to the adapter you want. Press Enter.

- c. From the **FAST!Util** menu, scroll down and select **Select Host Adapter**.
- d. Scroll up and highlight **Configuration Settings**. Press Enter.
- e. From the **Configuration Settings** menu, click **Host Adapter Settings**.
- f. Write down the 16-digit alphanumeric string that is displayed on your screen.

When you configure storage on the storage unit for fibre-channel attachment, you must specify the worldwide name (WWN) of the fibre-channel adapter in the IBM System Storage DS Storage Manager.

On Linux systems, the required WWN displays in the `/var/log/messages` system log file when you load the driver. This location, however, depends on the system configuration. You can perform this task automatically during startup, or you can manually perform the task to load the module using the `insmod` or `modprobe` command.

The following is example output for the `more /var/log/messages` command.

```
# more /var/log/messages
...
Nov 2 08:55:05 skywalker kernel: (scsi): Found a QLA2200 @ bus 1, device 0x1,
irq 20, iobase 0x2300
Nov 2 08:55:10 skywalker kernel: scsi(3): Configure NVRAM parameters...
Nov 2 08:55:10 skywalker kernel: scsi(3): Verifying loaded RISC code...
Nov 2 08:55:10 skywalker kernel: scsi(3): Verifying chip...
Nov 2 08:55:10 skywalker kernel: scsi(3): Waiting for LIP to complete...
Nov 2 08:55:10 skywalker kernel: scsi(3): LOOP UP detected
Nov 2 08:55:10 skywalker kernel: scsi3: Topology - (F_Port), Host Loop address 0xffff
Nov 2 08:55:10 skywalker kernel: scsi(3): Waiting for LIP to complete...
Nov 2 08:55:10 skywalker kernel: scsi3: Topology - (F_Port), Host Loop address 0xffff
Nov 2 08:55:10 skywalker kernel: scsi-qla1-adapter-node=200000e08b02e534;
Nov 2 08:55:10 skywalker kernel: scsi-qla1-adapter-port=210000e08b02e534;
Nov 2 08:55:10 skywalker kernel: scsi-qla1-target-0=5005076300c08b1f;
...
```

Managing SCSI disk connectivity

In order to effectively manage SCSI disk on an Intel host running Linux, review the review the information in this section.

This section contains the following SCSI disk connectivity detailed information:

- Overview of the Linux SCSI subsystem
- LUN identification for the Linux host system
- Dynamic SAN fabric reconfiguration
- LUN detection procedures
- Linux Device Naming
- SCSI Disk Limitations: Defining the number of disk devices on Linux
- SCSI disk problem identification and resolution

Linux SCSI subsystem overview

This section provides an overview of SCSI disk connectivity.

The Linux SCSI subsystem consists of three layers:

- The upper layer consists of specific device type drivers that are closest to user-space, such as the disk driver, tape driver, and CD-ROM driver.
- The lower layer consists of drivers, such as the QLogic and Emulex host bus adapter (HBA) drivers, which are closest to the hardware.
- The middle layer is the SCSI core that glues the upper and lower layer drivers together.

Depending on the kernel distribution, the drivers may be compiled into the kernel or built as modules which are loaded into the kernel. The sd driver is the SCSI disk driver, or block driver, which is named sd_mod when built as a module. The SCSI mid-layer driver is known as scsi_mod when built as a module.

Typically, in most distributions these drivers are built as modules and will be loaded at boot as part of the initrd images. If they are not loaded on boot, and are required at boot time, then an initrd image should be built to include them. For 2.4 kernels this is done by modifying the /etc/modules.conf file. For 2.6 kernels, this is done by modifying the /etc/modprobe.conf file and the /etc/sysconfig/kernel file. After modifying these files, the mkinitrd command needs to be run to pick up the changes.

To check if the drivers are built as modules and are currently loaded, check for sd_mod and scsi_mod in the lsmod command output. The following is example output:

```
# lsmod | grep sd
sd_mod                13440 296
scsi_mod              104068 6  [qla2300 st sr_mod sg ncr53c8xx sd_mod]
```

Note: If the SCSI mid-layer driver is compiled as a module, it is important that scsi_mod is loaded before the upper and lower level drivers and unloaded last.

LUN identification for the Linux host system

This section describes LUN identification for the Linux host system.

In Linux 2.4 kernels, SCSI devices are discovered by scanning the SCSI bus when the host adapter driver is loaded. A list of devices which have been discovered and are recognized by the SCSI subsystem are listed in /proc/scsi/scsi. It is good practice to use the cat command to display the output of /proc/scsi/scsi to verify that the correct number of LUNs has been recognized by the kernel. The following is example output:

```
# cat /proc/scsi/scsi
# cat /proc/scsi/scsi | less
Attached devices:
Host: scsi1 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM-PSG Model: DDYS-T36950M M Rev: S96E
  Type: Direct-Access ANSI SCSI revision: 03
Host: scsi1 Channel: 00 Id: 08 Lun: 00
  Vendor: IBM Model: YGLv3 S2 Rev: 0
  Type: Processor ANSI SCSI revision: 02
Host: scsi4 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM Model: 1750500 Rev: 5.17
  Type: Direct-Access ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM Model: 1750500 Rev: 5.17
  Type: Direct-Access ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 02
  Vendor: IBM Model: 1750500 Rev: 5.17
  Type: Direct-Access ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 03
  Vendor: IBM Model: 1750500 Rev: 5.17
  Type: Direct-Access ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 04
  Vendor: IBM Model: 1750500 Rev: 5.17
  Type: Direct-Access ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 05
  Vendor: IBM Model: 1750500 Rev: 5.17
  Type: Direct-Access ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 06
  Vendor: IBM Model: 1750500 Rev: 5.17
  Type: Direct-Access ANSI SCSI revision: 05
Host: scsi4 Channel: 00 Id: 00 Lun: 07
  Vendor: IBM Model: 1750500 Rev: 5.17
  Type: Direct-Access ANSI SCSI revision: 05
```

Note: The `/proc/scsi/scsi` list of devices is not dynamic and does not reflect state changes caused by fabric changes. Please see the section on “Dynamic SAN fabric reconfiguration” for more information on how to deal with fabric changes.

Starting with the Linux 2.6 kernels, the `/proc` filesystem is migrating over to the improved `/sys` filesystem. The `/sys` filesystem supports dynamic changes such as adding and removing LUNs without requiring a host adapter driver reload or rebooting the host system. Typically, more accurate information about what SCSI devices are recognized by the host system is attained by checking the contents of the directory `/sys/class/scsi_host/hostN`, where N is host adapter ID number. `lsscsi` is a user-space tool that uses the information in `/sys` to display a summarized list of all the recognized devices. This tool may not currently be available on all Linux 2.6 distribution kernels.

Note: The `lsscsi` command is only supported on the SLES distributions.

Dynamic SAN fabric reconfiguration

This section provides four methods that you can use to force the Linux operating system to recognize disk that are added or removed from the fabric.

When you add or remove disks to the fabric, you can use any of the following four ways to force the Linux host to recognize these changes:

1. Reboot the host
2. Unload and reload the host adapter driver
3. Rescan the bus by echoing the `/sys` filesystem (only for Linux 2.6 kernels)
4. Manually add and remove SCSI disks by echoing the `/proc` or `/sys` filesystem

Reboot the Host or Unload and Reload the host adapter driver

Since devices are discovered by scanning the SCSI bus, it is typically easiest to rescan the SCSI bus to detect any SAN fabric changes. A bus rescan is automatically triggered by reloading the host adapter driver or by rebooting the system.

Before unloading the host adapter driver or rebooting the host, you must:

1. Stop all I/O
2. Unmount all file systems
3. If SDD is being used, unload the SDD driver with the **sdd stop** command before reloading the host adapter driver. After the host adapter driver is reloaded then reload SDD with the **sdd start** command.

Reloading the host adapter driver assumes that the host adapter driver is built as a module. Rebooting the system works regardless of whether or not the host adapter driver is compiled into the kernel or as a module.

Rescan the bus by echoing the /sys filesystem (only for Linux 2.6 kernels)

For Linux 2.6 kernels only, a rescan can be triggered through the /sys interface without having to unload the host adapter driver or reboot the system. The following command will scan all channels, targets, and LUNs on host H.

```
echo "- - -" > /sys/class/scsi_host/hostH/scan
```

Manually add and remove SCSI disks

You can use the following commands to manually add and remove SCSI disk.

Note: In the following command examples, H, B, T, L, are the host, bus, target, and LUN IDs for the device.

You can unconfigure and remove an unused SCSI disk with the following command:

```
echo "scsi remove-single-device H B T L" > /proc/scsi/scsi
```

If the driver cannot be unloaded and loaded again, and you know the host, bus, target and LUN IDs for the new devices, you can add them through the /proc/scsi/scsi file using the following command:

```
echo "scsi add-single-device H B T L" > /proc/scsi/scsi
```

For Linux 2.6 kernels, devices can also be added and removed through the /sys filesystem. Use the following command to remove a disk from the kernel's recognition:

```
echo "1" > /sys/class/scsi_host/hostH/device/H:B:T:L/delete
```

or, as a possible variant on other 2.6 kernels, you can use the command:

```
echo "1" > /sys/class/scsi_host/hostH/device/targetH:B:T/H:B:T:L/delete
```

To reregister the disk with the kernel use the command:

```
echo "B T L" > /sys/class/scsi_host/hostH/scan
```

Note: The Linux kernel does not assign permanent names for the fabric devices in the /dev directory. Device file names are assigned in the order in which devices are discovered during the bus scanning. For example, a LUN might be /dev/sda. After a driver reload, the same LUN might become /dev/sdce. A fabric reconfiguration might also result in a shift in the host, bus, target and LUN IDs, which makes it unreliable to add specific devices through the /proc/scsi/scsi file.

LUN detection procedures

This section describes LUN detection procedures for the Linux host system.

If you have a Linux driver that does not automatically configure any LUNs other than LUN 0, you can manually configure the other LUNs, depending on the parameters and settings used for the SCSI mid-layer driver. Figure 9 shows an example of the /proc/scsi/scsi file for a Linux host that only configures the first LUN, LUN 0, on each host adapter port.

```
# cat /proc/scsi/scsi
...
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM-PSG   Model: DPSS-318350M  F  Rev: S9HA
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi0 Channel: 00 Id: 15 Lun: 00
  Vendor: IBM       Model: TP4.6 V41b3    Rev: 4.1b
  Type:   Processor                      ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM       Model: 2105800        Rev: .294
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM       Model: 2105800        Rev: .294
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM       Model: 2105800        Rev: .294
  Type:   Direct-Access                    ANSI SCSI revision: 03
```

Figure 9. Example of a /proc/scsi/scsi file from a Linux host that only configures LUN 0

There are two ways to workaround the issue of only having LUN 0 configured:

1. Create a script to manually add devices into /proc/scsi/scsi
2. Detect LUNs automatically at system boot by modifying the initial ram-disk (initrd)

Create a script to echo the /proc filesystem

Use the **scsi add-single-device** command to consecutively configure all of the LUNs that are assigned to your host system. Write a script that repeats the **scsi add-single-device** command for each LUN on each ID for each host adapter. The script must scan all host adapter ports and identify all of the LUNs that are assigned to each port.

After you run the script, you can view all of the assigned LUNs in the /proc/scsi/scsi file.

Figure 10 on page 112 shows an excerpt of an example /proc/scsi/scsi file for a Linux host after a script has configured every LUN.

```
# cat /proc/scsi/scsi
...
Host: scsi3 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 2105800      Rev: .294
  Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: 2105800      Rev: .294
  Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: 2105800      Rev: .294
  Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 02
  Vendor: IBM      Model: 2105800      Rev: .294
  Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 03
  Vendor: IBM      Model: 2105800      Rev: .294
  Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 04
  Vendor: IBM      Model: 2105800      Rev: .294
  Type:   Direct-Access                ANSI SCSI revision: 03
...
```

Figure 10. Example of a `/proc/scsi/scsi` file for a Linux host with configured LUNs

Detect LUNs automatically at system boot

The second method of configuring LUNs for a Linux system with only LUN 0 configured involves setting the parameter for the SCSI mid-layer driver that controls how many LUNs are scanned during a SCSI bus scan. The following procedure works for both 2.4 and 2.6 kernels, but it assumes the SCSI mid-layer driver is compiled as a `scsi_mod` module that is loaded automatically at system boot time. For Linux 2.4 kernels, to set the maximum number of disk devices under Linux to properly detect all volumes, you need to set the `max_scsi_luns` option for the SCSI mid-layer driver. For example, if `max_scsi_luns` is set to 1 this limits SCSI bus scans to only LUN 0. This value should be set to the respective maximum number of disks the kernel can support, for example, 128 or 256. In Linux 2.6 kernels, the same procedure applies, except that the parameter has been renamed from `max_scsi_luns` to `max_luns`.

1. Edit the `/etc/modules.conf` file.
2. Add the following line:

```
options scsi_mod max_scsi_luns=<n> (where <n> is the total number of
luns to probe.
```
3. Save the file.
4. Run the **mkinitrd** command to rebuild the ram-disk associated with the current kernel. You can use the following figures examples of what **mkinitrd** command to run for your operating system. `<kernel>` refers to the 'uname -r' output which displays the currently running kernel level, for example: 2.4.21-292-smp. For SUSE distributions, use the following command:

```
cd /boot
mkinitrd -k vmlinuz-<kernel> -i initrd-<kernel>
```

For Red Hat distributions, use the following command:

```
cd /boot
mkinitrd -v initrd-<kernel>.img <kernel>
```
5. Reboot the host.
6. Verify that the boot files are correctly configured for the newly created initrd image in the `/boot/grub/menu.lst` file.

Linux device naming

This section provides information to help you understand naming patterns for devices on Linux.

Special device files are used by kernel drivers to control a device. There can be more than one special device file that maps to the same physical device. For example, in a multipath environment with four paths to the same device, four different device files will map to the same physical device.

The device files are located in the `/dev` directory and are addressed by a major and minor number pair. Fibre-channel attached devices are managed as SCSI disk devices by the `sd` driver. Thus, each of the attached storage unit LUNs has a special device file in the Linux directory `/dev`.

SCSI disk devices have special device files that start with an 'sd' prefix, and have the following naming format:

```
/dev/sd[a-z][a-z][1-15]
```

Names without a trailing digit refer to the whole disk, while names with a trailing digit refer to a partition of that whole disk. By convention, SCSI disks have a maximum of 16 minor numbers mapped to a single disk. Thus, for each whole disk, there is a maximum of 15 partitions per disk because one minor number is used to describe the entire disk (for example `/dev/sda`), and the other 15 minor numbers are used to refer to partitions for that disk (for example `/dev/sda1`, `/dev/sda2`, etc). The following example displays the device files for the whole disk `/dev/sda`, which has a major number of 8 and a minor number of 0, and its 15 partitions.

```
# ls -l /dev/sda*
brw-rw---- 1 root disk 8, 0 May 24 08:09 /dev/sda
brw-rw---- 1 root disk 8, 1 May 24 08:09 /dev/sda1
brw-rw---- 1 root disk 8, 10 May 24 08:09 /dev/sda10
brw-rw---- 1 root disk 8, 11 May 24 08:09 /dev/sda11
brw-rw---- 1 root disk 8, 12 May 24 08:09 /dev/sda12
brw-rw---- 1 root disk 8, 13 May 24 08:09 /dev/sda13
brw-rw---- 1 root disk 8, 14 May 24 08:09 /dev/sda14
brw-rw---- 1 root disk 8, 15 May 24 08:09 /dev/sda15
brw-rw---- 1 root disk 8, 2 May 24 08:09 /dev/sda2
brw-rw---- 1 root disk 8, 3 May 24 08:09 /dev/sda3
brw-rw---- 1 root disk 8, 4 May 24 08:09 /dev/sda4
brw-rw---- 1 root disk 8, 5 May 24 08:09 /dev/sda5
brw-rw---- 1 root disk 8, 6 May 24 08:09 /dev/sda6
brw-rw---- 1 root disk 8, 7 May 24 08:09 /dev/sda7
brw-rw---- 1 root disk 8, 8 May 24 08:09 /dev/sda8
brw-rw---- 1 root disk 8, 9 May 24 08:09 /dev/sda9
```

For Red Hat, the kernel automatically creates the device files for 128 devices. For SUSE, there are only special device files for the first 16 disks. You must create the device files for additional disks by using the **mknod** command. For 2.6 kernels, the special device files are only created when they are detected and recognized by the kernel. The `/proc/partitions` file lists all the 'sd' devices that are recognized by the SCSI disk driver, including the sd name, major number, minor number, and size of each disk device.

The following is an example `/proc/partitions` file:


```
# cat /proc/partitions
major minor #blocks name

 8      0 17774160 sda
 8      1 1052226 sda1
 8      2 208845 sda2
 8      3 10490445 sda3
 8     16 976576 sdb
 8     32 976576 sdc
 8     48 976576 sdd
 8     64 976576 sde
 8     80 976576 sdf
 8     96 976576 sdg
 8    112 976576 sdh
 8    128 976576 sdi
 8    144 976576 sdj
 8    160 976576 sdk
 8    176 976576 sdl
 8    192 976576 sdm
 8    208 976576 sdn
 8    224 976576 sdo
 8    240 976576 sdp
65      0 976576 sdq
65     16 1048576 sdr
65     32 1048576 sds
65     48 1048576 sdt
65     64 1048576 sdu
65     80 1048576 sdv
65     96 1048576 sdw
65    112 1048576 sdx
65    128 1048576 sdy
65    144 1048576 sdz
65    160 1048576 sdaa
65    176 1048576 sdab
65    192 1048576 sdac
65    208 1048576 sdad
65    224 1048576 sdae
65    240 1048576 sdaf
66      0 1048576 sdag
66     16 1048576 sdah
66     32 1048576 sdai
66     48 1048576 sdaj
66     64 1048576 sdak
66     80 1048576 sdal
66     96 1048576 sdam
66    112 1048576 sdan
66    128 1048576 sdao
66    144 1048576 sdap
66    160 1048576 sdaq
66    176 1048576 sdar
66    192 1048576 sdas
66    208 1048576 sdat
66    224 1048576 sdau
66    240 1048576 sdav
```

SCSI disk limitations: Defining the number of disk devices on Linux

This section provides information to help you understand SCSI disk limitations.

The Linux kernel uses a static major and minor number address mechanism. There are a limited number of major numbers that the kernel reserves for SCSI disk devices. Thus, there are limitations to the number of SCSI disk devices that can exist depending on the major numbers available. The number of major numbers made available for SCSI disk devices varies depending on the distribution kernel level.

For Linux 2.4 kernels, the maximum number of configurable devices that are available on supported Linux host systems is either 128, 256, or 2304. Eight bits are reserved for minor numbers, thus there are 28, or 256, minor numbers available for each major number. By convention, for SCSI disks 16 minor numbers are reserved for each major number, thus there are up to 15 allowable partitions per disk. With 16 minor numbers used to describe each disk and its partitions, and 256 possible minor numbers, then there are $256 \div 16 = 16$ possible disks that can be described by a single major number.

Traditionally, there have been eight major numbers reserved for SCSI devices. Those major numbers are 8, 65, 66, 67, 68, 79, 70, and 71. RHEL 2.1 and RHEL3 use the 8 reserved major numbers. Thus, with 8 reserved major numbers, and 16 possible disks per major number, these kernels can describe up to $8 \times 16 = 128$ possible disks. The earlier SuSE SLES8 kernels can support up to 256 devices because they reserve 16 major numbers. The additional major numbers are 128, 129, 130, 131, 132, 133, 134, and 135.

Newer kernels are moving to a dynamic configuration method that will use any available major numbers after the first 16 reserved major numbers have been used. The latest SuSE SLES 8 kernel errata includes this change and imposes a limit of no more than 144 major numbers allocated to SCSI disk devices. Thus these kernels can support up to 2304 disk devices.

In general, the following formula provides the maximum number of devices for Linux host systems: Number of devices = (number of major numbers) \times (number of minor numbers) \div (number of partitions). For example, Number of devices = $8 \times 256 \div 16 = 128$ or Number of devices = $144 \times 256 \div 16 = 2304$.

For the Linux 2.6 kernels the number of major and minor bits has been increased to 12 and 20 bits respectively, thus Linux 2.6 kernels can support thousands of disks. There is still the limitation of only up to 15 partitions per disk.

Other factors that limit the number of disk devices

If you load a host adapter driver as a module in Linux, there is a kernel limit to the total number of disks that you can configure. This limit might be less than the total supported by the kernel (usually 128 or 256). The first module that is loaded on the system might configure discovered disks up to the maximum number of disks supported by the kernel. Subsequent drivers are limited to configuring a fixed, smaller number of drives. These drivers all share a pool of device structures that are statically allocated after the first host adapter driver is loaded. The number of allocated device structures is set with the `CONFIG_SD_EXTRA_DEVS` parameter when the kernel compiles. Traditionally, this value has been 40, but newer Red Hat kernels have it set to 128 or 256.

For example, consider a system with an internal SCSI Adaptec host adapter and one disk. The system contains fibre connected disks and a `CONFIG_SD_EXTRA_DEVS` parameter set to 40. The Adaptec host adapter driver loads first and configures the single attached disk. At this point, the system is limited to a total of 41 SCSI disks instead of the full 128 or 256. If there are more than 40 disks visible in the fibre channel fabric, the system configures and uses only the first 40. All of the fabric devices might be listed in the `/proc/scsi/scsi` directory, but only the first 40 are configured with the SCSI disk driver, `scsi_mod`. You are more likely to encounter this constraint in multipath fabric situations.

For the later SuSE SLES8 kernels that support 2304 SCSI disks, the CONFIG_SD_EXTRA_DEVS parameter was replaced with the CONFIG_SD_MAX_MAJORS parameter with default value of 144. For Linux 2.6 kernels, both these parameters are obsolete and are no longer needed.

SCSI disk problem identification and resolution

This section describes a few problems that might occur on your storage unit when you are using SCSI disks.

- In response to errors in the SAN, the kernel might permanently disable a LUN and log a message stating "device set offline" and the specific device. If this happens on the 2.4 kernel, there is no way to bring the LUN online except for unloading the low-level device driver and reloading the driver or rebooting the system.
- On 2.6 kernels, the device can be brought back online using one of the following methods:
 - Redhat: `echo "running" >/sys/class/scsi_host/hostH/device/targetH:C:T/H:C:T:L/state`
 - SLES: `echo "1" > /sys/class/scsi_host/hostH/device/H:C:T:L /online`
- The system might periodically list processes in the D-state (see the `ps` command help page), which corresponds to an uninterruptible process, due to the process waiting in the kernel. In error situations, a process might become permanently stuck in this state, and require a system reboot to recover.
- The Linux kernel buffer cache is designed to discard dirty buffers after an input/output (I/O) error when the system memory resources are constrained. An application that is attempting to use the `fsync()` command to verify that its writes have completed successfully will receive an indication of success from the command once the writes successfully complete. Some kernels have a bug in the `kswapd` daemon, that makes it likely that the system will perceive itself to be in a state of constrained memory. Multipathing can reduce the risk of this silent data loss by providing a means to retry failed I/O operations and hide the failure from the buffer cache.

Creating a DS6000 file system and partitions for an Intel host running Linux

To attach an Intel host running the Linux operating system to a DS6000 you need to create a Linux file system on the DS6000. You can also partition the DS6000 disk for use by your Linux host by using the `fdisk` utility.

If you choose to partition the storage space on the DS6000, you must use the `fdisk` utility to partition the disk before you create a file system.

You have to specify the special device file of the disk you want to partition when executing `fdisk`. The following is an example of the different options for the `fdisk` utility.

```
# fdisk /dev/sdb

Command (m for help): m
Command action
a    toggle a bootable flag
b    edit bsd disklabel
c    toggle the dos compatibility flag
d    delete a partition
l    list known partition types
m    print this menu
n    add a new partition
o    create a new empty DOS partition table
p    print the partition table
q    quit without saving changes
s    create a new empty Sun disklabel
t    change a partitions system id
u    change display/entry units
v    verify the partition table
w    write table to disk and exit
x    extra functionality (experts only)
```

The following example shows how to use the **n** option of the **fdisk** utility to create a primary partition on the disk **/dev/sdb**.

```
Command (m for help): n

Command action
e    extended
p    primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-953, default 1): Enter
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-953, default 953): Enter
Using default value 953

Command (m for help): p

Disk /dev/sdb: 64 heads, 32 sectors, 953 cylinders
Units = cylinders of 2048 * 512 bytes

Device Boot Start End Blocks Id System
/dev/sdb1 1 953 975856 83 Linux
```

Assigning the system ID to the partition of an Intel or AMD host running Linux

This section provides the steps you need to assign the system ID to a Linux partition using the **fdisk** command.

1. Assign the system partition ID.
2. Write the information to the partition table on the disk.
3. Exit the **fdisk** program.

Figure 11 on page 118 shows the assignment of the Linux system ID to the partition (hex code 83).

```

Command (m for help): t
Partition number (1-4): 1

Hex code (type L to list codes): 83

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
SCSI device sdb: hdwr sector= 512 bytes. Sectors= 1953152 [953 MB] [1.0 GB]
sdb: sdb1
SCSI device sdb: hdwr sector= 512 bytes. Sectors= 1953152 [953 MB] [1.0 GB]
sdb: sdb1

WARNING: If you have created or modified any DOS 6.x partitions, please see the
fdisk manual page for additional information.
Syncing disks.
[root@yahoo /data]#

```

Figure 11. Example of assigning a Linux system ID to the partition

Creating a file system for an Intel or AMD host running Linux

This section describes how you create and use file systems on the storage unit for an Intel or ADM host running the Linux operating system.

After you partition the disk as described in “Creating a DS6000 file system and partitions for an Intel host running Linux” on page 116, the next step is to create a file system. Figure 12 shows an example of the EXT2 Linux file system (which is nonjournaled) by using the **mke2fs** command.

```

[root@yahoo /data]# mke2fs /dev/sdb1
mke2fs 1.18, 11-Nov-1999 for EXT2 FS 0.5b, 95/08/09
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
122112 inodes, 243964 blocks
12198 blocks (5.00%) reserved for the super user
First data block=0
8 block groups
32768 blocks per group, 32768 fragments per group
15264 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376

Writing inode tables: done
Writing superblocks and filesystem accounting information: done
[root@yahoo /data]#

```

Figure 12. Example of creating a file with the mke2fs command

Figure 13 on page 119 shows an example of the EXT2 Linux file system (which is nonjournaled) by using the **mkfs** command.

```

[root@yahoo /data]# mkfs -t ext2 /dev/sdb1
mke2fs 1.18, 11-Nov-1999 for EXT2 FS 0.5b, 95/08/09
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
122112 inodes, 243964 blocks
12198 blocks (5.00%) reserved for the super user
First data block=0
8 block groups
32768 blocks per group, 32768 fragments per group
15264 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376
Writing inode tables: done
Writing superblocks and filesystem accounting information: done
[root@yahoo /data]#

```

Figure 13. Example of creating a file with the mkfs command

Fibre-channel boot on an Intel or AMD host running Linux

The topics in this section contain support information for fibre-channel boot.

For additional fibre-channel boot support information, see the IBM System Storage Multipath Subsystem Device Driver User's Guide at <http://www.ibm.com/systems/support/storage/software/sdd>.

Creating a modules disk for SUSE Linux Enterprise Server 9.0

Complete this task to create a modules disk for SUSE Linux Enterprise Server 9.0.

1. Boot the server using the SLES 9 SP1 CD.
2. When you are prompted, select **Rescue System**.

Note: If you are using the x360/x440 servers, append "apic" to the kernel boot options string at the bottom of the menu screen.

3. Insert the SLES 9 CD 1 CD when you are prompted.
4. After the system boots, the Rescue login prompt appears. Log in as "root" and do not enter a password.
5. Remove the SLES 9 CD 1 CD, and insert the SLES 9 SP1 CD.
6. If you are using IDE CD-ROM drives, type `mount /dev/hda /cdrom` at the command line. If you are using SCSI, type `/dev/sda` instead of `hda`, or `/dev/sdb`, and so forth.
7. Change to the CD-ROM boot directory by typing `cd /cdrom/boot`.
8. Insert a diskette into the diskette drive. Create the modules disk by typing `dd if=modules2 of=/dev/fd0u1440`. When this is completed, a message is displayed that says "2880+0 records in 2880+0 records out".
9. Reboot with the SLES 9 SP1 CD in the CD-ROM drive.

Installing Linux over the SAN without an IBM subsystem device driver

Use this process to install Linux over the SAN without an IBM subsystem device driver (SDD). These instructions only apply to QLogic adapter cards or to the eServer BladeCenter HS20. These instructions do not apply to a server with Emulex adapter cards.

1. Boot the server.
2. Press the <CTL-Q> message for the Fast!Util when it appears.

3. In the QLogic FAST!Util, select the controller that has your remote boot disks attached.
4. Click **Configuration Settings** → **Host Adapter Settings** and change the BIOS enabled setting to **Enabled**.
5. Click **Configuration Settings** → **Selectable Boot Settings**.
 - a. Change the Selectable boot setting to **Enabled**.
 - b. Set **(Primary) Boot Port Name,Lun:** to the first device in your list.

Note: If you have assigned multiple LUNs and are using any of them to boot from (that is, if you have file systems such as /boot, /, /usr, /var, or the swap partition), you must list the LUNs in the subsequent entries, as well.

6. Save your changes.
7. Back out of the FAST!Util and reboot.
8. Insert modules disk modules2 into the diskette drive (see the previous procedure for creating modules disks).
9. Put the SP1 CD in the CD-ROM drive and reboot.
10. Select **Manual Installation**.

Note: If you are using a Summit machine, you must add “apic” to the boot options at the bottom of the menu screen.

11. Select the appropriate keyboard and press Enter.
12. After manual installation, select the **Kernel modules (hardware drivers)** and **Load IDE/RAID/SCSI modules** options.
13. Select the QLA23xxx module from the list, press Enter twice, leaving the parameters line blank.
14. Wait for the module to return from loading, and then press Enter.
15. Press the right-arrow key once to select the **Back** button, and press Enter again.
16. Select **Start Installation/Update** → **Install** → **CD-ROM**.
17. When you are prompted with “Please make sure that CD number 1 is in your drive!”, place the SLES 9 Install CD 1 CD into the CD-ROM drive and press **Enter**.
18. When you are prompted with “SUSE’s End User License for SLES”, select **Accept**, and then select the appropriate keyboard. Click **Accept**.
19. During detection, if you are asked whether you want to make a New Installation, select **New Installation** and continue.
20. Create the partitions. (Normally, there are three: one of approximately 200 MB for /boot, a second of approximately 2 GB for swap, and the rest for /.)
21. Make the package selections, making sure to include the kernel source package.
22. Check your boot device parameters. In most cases, add “apic=off” to the load string. If you have loaded the apic boot option at the outset, it is listed automatically. If you used an APIC-enabled installation and the apic boot option is not listed automatically, add it. Failure to add the apic boot option results in a panic condition for the server at reboot.
23. Check your time zone parameter and change it to the appropriate time zone.
24. Check the clock setting. If it defaults to “UTC”, change it to “Local”.
25. Click **Next** and then click **Yes** when you are prompted. The system formats the remote disk, and then installs the packages.

26. Switch CDs when you are prompted.

Note: To boot Linux, follow the instructions above to allow the server to boot.

Updating a module without an IBM subsystem device driver

Complete this task to update to a more recent module without IBM subsystem device driver (SDD).

1. You must boot into SUSE SLES 9 before you install the QLogic driver source and compile the necessary kernel and driver.
2. Download the current QLogic adapter driver. For additional information, see “Downloading and installing a host adapter driver” on page 30.
3. Type `cat /etc/sysconfig/kernel` at the command prompt and make sure that the QLA23xx driver is listed.
4. Type `mk_initrd` to pick up the new driver in the ramdisk image for this kernel, and then reboot.

Notes:

- a. If you are using lilo, you must rerun lilo to install the changed boot parameters into the MBR. To do this, type `mk_initrd && lilo` from the command line. GRUB users do not need to make any changes after running `mk_initrd`.
- b. If you ever need to upgrade the kernel, repeat this procedure for updating the module before you reboot. Make sure the `mk_initrd` command runs successfully before you reboot.

Installing Linux over the SAN with an IBM subsystem device driver

Use this process to install Linux over the SAN with an IBM subsystem device driver (SDD).

This procedure assumes that you have correctly configured the bootloader to boot from the single-pathed SAN device. It also assumes that the SDD rpm is installed on your system. This procedure describes how to copy the SDD files into the initial ramdisk (`initrd`) and edit the `linuxrc` script, which is processed when the kernel mounts the initial ramdisk at boot time. For current information about the SDD rpm, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide*, at <http://www.ibm.com/systems/support/storage/software/sdd>.

Perform the following steps to install Red Hat and SUSE with an IBM SDD:

1. Make a backup of the existing `_initrd_file_` by typing the following commands:

```
cd /boot
cp _initrd_file_ _initrd_file.bak
```

2. Uncompress the image by typing the following command:

```
zcat _initrd_file_.bak > _initrd_file_
```

3. Set up a loopback device for the image by typing the following command:

```
losetup /dev/loop0 /path/to/your/_initrd_file_
```

4. Fix any errors that might exist on the file system by typing the following command:


```
e2fsck -f /dev/loop0
```

5. Determine the size of the initrd file system by typing the following command:

```
df /dev/loop0
```

Ensure that you have sufficient space in the /boot directory (or other home of your initrd files) to store considerably larger initrd files (for example, files with a size of 32 MB each). If there is not sufficient space in your /boot directory, you can perform the following steps in a temporary directory and then copy the compressed initrd file (a few megabytes instead of 32 MB) back into the /boot directory.

If the file system is not 32 MB or larger or does not have much free space, you must enlarge it to 32 MB by typing the following command:

```
losetup -d /dev/loop0  
dd if=/dev/zero of=_initrd_file_ seek=33554432 count=1bs=1
```

Note: If the file is already 32 MB or larger, do not perform this step because it is unneeded and can corrupt the initial ramdisk file system.

On SUSE, you might need to create an even larger initial ramdisk (for example, 48 MB or 48×1024×1024 would result in a seek=50331648). If the initrd file is sufficiently large, skip ahead to mount the loopback device. See step 9.

6. Set up a loopback device for the image by typing the following command:

```
losetup /dev/loop0 /path/to/your/_initrd_file_
```

7. Ensure that you have a clean file system by again typing the following command:

```
e2fsck -f /dev/loop0
```

If you still have errors on the file system, the previous dd step was not performed correctly and it corrupted the initrd. You now must delete the loopback device by typing `losetup -d /dev/loop0` and restart the procedure from the beginning.

8. Resize the file system by typing the following command:

```
resize2fs /dev/loop0
```

Note: Resizing automatically expands the file system so that it uses all available space.

9. Mount the loopback device by typing the following command:

```
mount /dev/loop0 /mnt/tmp
```

10. You now have an initrd file system that is 32 MB. You can add additional files by typing the following command:

```
cd /mnt/tmp
```


11. If you have not already added your host adapter driver to the initrd file using the standard `mk_initrd` or `mkinitrd` process (depending on your distribution), you must manually copy the module files for the host adapter driver. You also must manually copy the SCSI core and SCSI disk drivers into the initrd file system and add the appropriate `insmod` command to the `linuxrc` script.
12. On SUSE, you must create the `etc/`, `proc/`, and `sysroot/` directories in the initrd file system.
You can also add `echo` commands into the `linuxrc` script after the host adapter load and mounting `/proc` to force the addition of LUNs through `/proc/scsi/scsi` if the device discovery is not occurring automatically.
13. Create the SDD directories in the initrd file system by typing the following commands:

```
mkdir -p opt/IBMsdd/bin
chmod -R 640 opt/IBMsdd/
```

14. For the SDD, you must copy the following files to the initrd file system.

Note: Ensure that you are in the `/mnt/tmp` directory when you perform the copies.

File Names	Target Location
<code>/etc/vpath.conf</code>	<code>etc/</code>
<code>/etc/group</code>	<code>etc/</code>
<code>/etc/passwd</code>	<code>etc/</code>
<code>/etc/nsswitch.conf</code>	<code>etc/</code>
<code>/opt/IBMsdd/sdd-mod.o-CORRECT_VERSION</code>	<code>lib/sdd-mod.o</code>
<code>/opt/IBMsdd/bin/*</code>	<code>opt/IBMsdd/bin/</code>
<code>/lib/libc.so.6</code>	<code>lib/</code>
<code>/lib/ld-linux.so.2</code>	<code>lib/</code>
<code>/lib/libacl.so.1</code>	<code>lib/</code>
<code>/lib/libattr.so.1</code>	<code>lib/</code>
<code>/lib/libdl.so.2</code>	<code>lib/</code>
<code>/lib/libm.so.6</code>	<code>lib/</code>
<code>/lib/libpthread.so.0</code>	<code>lib/</code>
<code>/lib/libnss_files.so.2</code>	<code>lib/</code>
<code>/lib/librt.so.1</code>	<code>lib/</code>
<code>/bin/awk, chmod, chown, cp, date, grep, ls, mknod, mount, ps, rm, sed, sh, tar, unmount</code>	<code>bin/</code>
<code>/dev/sd[a-z], sd[a-z][a-z]</code> For example, <code>tar cps /dev/sd[a-z] /dev/sd[a-z][a-z] tar xps</code>	<code>dev/</code>

15. For Red Hat, you must copy the following additional files to the file system:

File Names	Target Location
<code>/lib/libproc.so.2.0.7</code>	<code>lib/</code>
<code>/lib/libpcre.so.0</code>	<code>lib/</code>

File Names	Target Location
/lib/libtermcap.so.2	lib/
/bin/ash.static	bin/ash

16. For SUSE, you must copy the following additional files to the file system:

File Names	Target Location
/lib/libreadline.so.4	lib/
/lib/libhistory.so.4	lib/
/lib/libncurses.so.5	lib/
etc/nsswitch.conf Note: The etc/nsswitch.conf file must have its password and group entries changed to point to files instead of compat.	N/A

17. The following changes must be make to the initrd linuxrc script:

- a. For Red Hat, remove the following block of commands from the end of the file:

```
echo Creating block devices
mkdevices /dev
echo Creating root device
mkroot dev /dev/root
echo 0x0100 > /proc/sys/kernel/real-root-dev
echo Mounting root filesystem
mount -o defaults -ro - -t ext2 /dev/root /sysroot
pivot_root /sysroot /sysroot/initrd
umount /initrd/proc
```

You must change the first line of the linuxrc script to invoke the ash shell instead of the nash shell.

- b. If the /proc file system is not already explicitly mounted in the linuxrc script, append the following mount command:

```
mount -n -tproc /proc /proc
```

- c. To configure SDD, append the following commands to the end of the linuxrc script:

```
insmod /lib/sdd-mod.o
/opt/IBMsdd/bin/cfgvpath
```

Mount the systems root file system so that you can copy configuration information to it. For example, if you have an ext3 root file system on /dev/vpatha3, enter /bin/mount -o rw -t ext3, or /dev/vpatha3 /sysroot, or for a reiserfs root file system on /dev/vpatha3, enter /bin/mount -o rw -t reiserfs /dev/vpatha3 /sysroot

To copy the dynamically created device special files onto the systems root file system, enter the following commands:

```
tar cps /dev/IBMsdd /dev/vpath* | (cd /sysroot && tar xps)
/bin/umount /sysroot
```

You must define the root file system to the kernel. Traditionally, this information is passed to the bootloader as a string, for example /dev/vpatha3 and translated to a hexadecimal representation of the major

and the minor numbers of the device. If the major and minor numbers equals 254,3, these numbers are represented in hex as 0xFE03. The linuxrc script passes the hexadecimal value into /proc with the following commands:

```
echo 0xFE03 > /proc/sys/kernel/real-root-dev  
/bin/umount /proc
```

18. Edit the system fstab and change all the system mount points from LABEL or /dev/sd mount sources to their equivalent /dev/vpath. Refer to step 23 for dangers of booting by label in a multipath configuration.
19. Copy the system fstab to the initrd etc/ directory.
20. Unmount the image and remove the loopback binding by typing the following commands:

```
umount /mnt/tmp  
losetup -d /dev/loop0
```

21. Compress the image by typing the following commands:

```
gzip -9 _initrd_file_  
mv _initrd_file_.gz _initrd_file_
```

22. Append the following to your boot parameters (for example, in lilo.conf, grub.conf, or menu.lst):

```
ramdisk_size=34000
```

If you created a larger initrd file system, make this value large enough to cover the size.

23. For completeness, change the bootloader append for the root parameter of the kernel to the appropriate vpath device. For example, root=/dev/vpatha5. However, the previous steps override this value by passing the corresponding hex major, minor into the /proc file system within the initrd linuxrc script.

Note: If you boot by LABEL, there is a risk that the first device that is found in the fabric with the correct label could be the wrong device or that it is an sd single-pathed device instead of a vpath multipathed device.

24. Reboot the server. It will boot with the root file system on a vpath device instead of on an sd device.

Figure 14 on page 126 and Figure 15 on page 126 illustrate a complete linuxrc file for Red Hat and for SUSE.

```
#!/bin/ash

echo "Loading scsi_mod.o module"
insmod /lib/scsi_mod.o
echo "Loading sd_mod.o module"
insmod /lib/sd_mod.o
echo "Loading qla2310.o module"
insmod /lib/qla2310.o
echo "Loading jbd.o module"
insmod /lib/jbd.o
echo "Loading ext3.o module"
insmod /lib/ext3.o
echo Mounting /proc filesystem
mount -t proc /proc /proc
insmod /lib/sdd-mod.o
/opt/IBMsdd/bin/cfgvpath
/bin/mount -o rw -t ext3 /dev/vpatha3 /sysroot
tar cps /dev/IBMsdd /dev/vpath* | (cd /sysroot && tar xps)
/bin/umount /sysroot
echo 0xFE03 > /proc/sys/kernel/real-root-dev
/bin/umount /proc
```

Figure 14. Example of a complete linuxrc file for Red Hat

```
#!/bin/ash

export PATH=/sbin:/bin:/usr/bin

# check for SCSI parameters in /proc/cmdline
mount -n -tproc none /proc
for p in `cat /proc/cmdline` ; do
    case $p in
        scsi*|*_scsi_*|llun_blkfst=*|max_report_luns=*)
            extra_scsi_params="$extra_scsi_params $p"
            ;;
    esac
done
umount -n /proc

echo "Loading kernel/drivers/scsi/scsi_mod.o $extra_scsi_params"
insmod /lib/modules/2.4.21-190-smp/kernel/drivers/scsi/scsi_mod.o $extra_scsi_params

echo "Loading kernel/drivers/scsi/sd_mod.o"
insmod /lib/modules/2.4.21-190-smp/kernel/drivers/scsi/sd_mod.o

echo "Loading kernel/drivers/scsi/qla2310_conf.o"
insmod /lib/modules/2.4.21-190-smp/kernel/drivers/scsi/qla2310_conf.o

echo "Loading kernel/drivers/scsi/qla2310.o"
insmod /lib/modules/2.4.21-190-smp/kernel/drivers/scsi/qla2310.o

echo "Loading kernel/drivers/scsi/aic7xxx/aic7xxx.o"
insmod /lib/modules/2.4.21-190-smp/kernel/drivers/scsi/aic7xxx/aic7xxx.o

echo "Loading kernel/fs/reiserfs/reiserfs.o"
insmod /lib/modules/2.4.21-190-smp/kernel/fs/reiserfs/reiserfs.o
mount -t proc /proc /proc
insmod /lib/sdd-mod.o
/opt/IBMsdd/bin/cfgvpath
/bin/mount -o rw -t reiserfs /dev/vpatha3 /sysroot
tar cps /dev/IBMsdd /dev/vpath* | (cd /sysroot && tar xps)
/bin/umount /sysroot
echo 0xFE03 > /proc/sys/kernel/real-root-dev
/bin/umount /proc
```

Figure 15. Example of a complete linuxrc file for SUSE

Intel or AMD running VMware ESX server host attachment

This chapter describes how you attach an Intel or Advanced Micro Devices (AMD) host system running VMware ESX Server to a DS6000 with fibre-channel adapters.

Note: The information in this chapter refers to VMware ESX Server v2.x and v3.x.

This chapter contains the following sections:

- Installing the Emulex adapter
- Installing the QLogic adapter and driver
- Defining the number of disk devices
- SCSI disk considerations
- Configuring an Intel or AMD host running VMware ESX Server
- Copy Services considerations

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>
- The VMware product documentation for the version of VMware that you are using. VMware documentation can be found at <http://www.vmware.com>.
- The LUN limitations according to the following information.

Note: The calculated limitation is a combination of the following VMware limitations:

- Maximum of 64 LUNs running I/Os concurrently and
- Less than 1024 LUNs managed by multipathing

Number of paths ¹	Maximum number of LUNs configured on storage / HBA
1 through 16	64
32	31
Note: 1. There should be either more host adapters on the host or more ports on storage.	

Installing the Emulex adapter for an Intel or AMD host running VMware ESX Server

This section describes how you install the Emulex adapter for an Intel or AMD host running VMware ESX Server.

Note: The LP9402DC adapter card is a 3.3 volt card only and requires a 3.3 volt PCI slot.

Single and dual-port fibre-channel interfaces with an Emulex adapter card support the following public and private loop modes:

- Target
- Public initiator

- Private initiator
- Target and public initiator
- Target and private initiator

To install the Emulex adapter, perform the following steps:

1. Record the Institute of Electrical and Electronics Engineers (IEEE) number that is printed on the card. You can use the IEEE number to determine the WWPN.

Note: You can find the WWPN in your `/proc/scsi/adapter_type/n` directory, where `adapter_type` is the host adapter type and `n` is the host adapter number for your card. Your host adapter type can be either `qlaxxxx` for QLogic adapters or `lpfccd` for Emulex adapters.

2. Refer to the installation instructions provided by your host adapter vendor for their specific installation instructions. See <http://www.emulex.com/ts/docoem/ibm/index.htm> for the latest documentation.

Installing the QLogic adapter and driver for an Intel or AMD host running VMware ESX Server

This section describes how you install and configure a QLogic adapter and driver for an Intel or AMD host system running VMware ESX Server.

Single- and dual-port fibre-channel interfaces with the QLogic adapters support the following public and private loop modes:

- Target
- Public initiator
- Private initiator
- Target and public initiator
- Target and private initiator

The following steps are an example configuration for a QLogic QLA23xx adapter. The configuration for your adapter might differ.

1. Install the QLogic adapter in the host system.
2. Connect the cable to the adapter and to the DS6000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS6000, see Host configuration using the DS Storage Manager.
3. Restart the server.
4. Press Ctrl+Q to get to the FAST!Util menu panel.
5. Use FAST!Util to set the parameters to the values that are listed in Table 7:

Table 7. Recommended settings for the QLogic adapter for an Intel or AMD host running VMware ESX Server

Parameters	Recommended settings
Basic Input/Output System (BIOS)	Disabled
Frame size	2048
Loop reset delay	5 (minimum)
Adapter hard loop ID	Disabled
Hard loop ID	0
Execution throttle	100

Table 7. Recommended settings for the QLogic adapter for an Intel or AMD host running VMware ESX Server (continued)

Parameters	Recommended settings
Fast command posting	Enabled
>4 GB addressing	Disabled (for 32-bit systems)
LUNs per target	0 or 128
Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes
Login retry count	<ul style="list-style-type: none"> • 60 if you are not running the SDD • 8 if you are running the SDD
Port down retry count	<ul style="list-style-type: none"> • 60 if you are not running the SDD • 8 if you are running the SDD
Driver load RISC code	Enabled
Enable database updates	No
Disable database load	No
IOCB allocation	256
Extended error logging	Disabled (You might want to enable this parameter if you use debugging.)
Type of connection	Point-to-point

6. Install and configure the adapter driver on the host. See the installation instructions that are provided by your host adapter vendor for specific adapter and adapter driver instructions. Go to “Downloading and installing a host adapter driver” on page 30 or http://support.qlogic.com/support/oem_ibm.asp.
7. Restart the server.

Defining the number of disk devices on a VMware ESX server

This section provides information for defining the number of disk devices on VMware ESX server.

The maximum number of devices that are supported on a VMware ESX server is 128 (256 for v3.x). However, the VMkernel supports up to 64 (128 for v3.x) concurrent I/Os. The embedded multipathing capability of the VMkernel can manage up to 1024 disks.

The maximum number of LUNs that can be used must accommodate the 3 previous limitations.

SCSI disk considerations for an Intel or AMD host running VMware ESX server

This section describes how you identify LUNs and disk devices, and provides information on persistent binding.

LUN identification for the VMware ESX console

This section provides LUN identification information for the VMware ESX console.

Each of the attached storage unit LUNs has a special device file in the Console OS directory /dev. There is a maximum of 128 (256 for v3.x), fibre-channel disks entries with the following format /dev/vsd [0 -127] (dev/sd nn for v3.x, where n is either a letter or a space).

From the VMware ESX console, the storage unit can be traced through several layers:

- Host adapter level: The host adapter driver provides the list of LUNs that are detected for each fibre-channel port.
- Vmhba level: VMware virtualizes the host adapter as a *vmhba* (VMware host bus adapter) from which you can retrieve the information that is related to the multipathing. In v3.x, you can use the `esxcfg-mpath` command to retrieve LUN and path data.

Host adapter level

The host adapter driver (`lpfcdd` for Emulex or `qla23xx` for QLogic) provides information on the port base. The content of the information varies according to the driver. The directory `/proc/scsi/ <driver name> / <adapter number>` provides the information for a specific adapter/port. Figure 16 on page 131 provides an example of QLogic output and Figure 17 on page 131 provides an example of Emulex output.


```

[root@denethor root]# cat /proc/scsi/qla2310/2
QLogic PCI to Fibre Channel Host Adapter for ISP23xx:
    Firmware version: 3.01.18, Driver version 6.04.02
Entry address = 0x4d8000
HBA: QLA2312 , Serial# H88769
Request Queue = 0x210ec000, Response Queue = 0x21100000
Request Queue count= 128, Response Queue count= 512
Total number of active commands = 0
Total number of interrupts = 4008395
Total number of IOCBs (used/max) = (0/600)
Total number of queued commands = 0
    Device queue depth = 0x10
Number of free request entries = 128
Number of mailbox timeouts = 0
Number of ISP aborts = 90
Number of loop resyncs = 300
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= <DOWN>, flags= 0x860813
Dpc flags = 0x0
MBX flags = 0x0
SRB Free Count = 4096
Port down retry = 30
Login retry count = 30
Commands retried with dropped frame(s) =0

SCSI Device Information:
scsi-qlal-adapter-node=200000e08b0c2109;
scsi-qlal-adapter-port=210000e08b0c2109;
scsi-qlal-target-0=5005076300c6a3bb;
scsi-qlal-target-1=200500a0b80f0a15;
scsi-qlal-target-2=5005076300c6a6b8;
scsi-qlal-target-3=5005076300c2a6b8;
scsi-qlal-target-4=5005076300cfa6b8;

SCSI LUN Information:
(Id:Lun)
( 4: 0): Total reqs 7, Pending reqs 0, flags 0x0, 1:0:81,
( 4: 1): Total reqs 7132, Pending reqs 0, flags 0x0, 1:0:81,
( 4: 2): Total reqs 7414, Pending reqs 0, flags 0x0, 1:0:81,
( 4: 3): Total reqs 2948, Pending reqs 0, flags 0x0, 1:0:81,
...

```

Figure 16. Example of QLogic output for a VMware ESX server

```

HBA: Emulex LightPulse LP9002 on PCI bus 06 device 28 irq 49
Emulex LightPulse LPFC Driver Version: 4.20q
SerialNum: MI23710129
Firmware Version: 3.90A7
Hdw: 2002606d
VendorId: 0xf90010df
Portname: 10:00:00:00:c9:2e:78:cc   Nodename: 20:00:00:00:c9:2e:78:cc

Link Up - Ready:
    PortID 0x71000a
    Fabric
    Current speed 2G

lpfc1t00 DID 71000e WWPN 50:05:07:63:00:cf:a6:b8 WWNN 50:05:07:63:00:c0:a6:b8

```

Figure 17. Example of Emulex output for a VMware ESX server

Vmhba level

The /proc/vmware/scsi directory presents 1 directory vmhba<N>, where <N> is the host adapter. The directory contains 1 entry per device (with a format <scsi id> : <LUN number>). Figure 18 on page 132 provides an example of a Vmhba

directory.

```
[root@denethor root]# ls /proc/vmware/scsi/vmhba0/
4:0  4:14 4:2  4:25 4:30 4:36 4:41 4:47 4:52 4:58 4:63 4:69 stats
4:1  4:15 4:20 4:26 4:31 4:37 4:42 4:48 4:53 4:59 4:64 4:7
4:10 4:16 4:21 4:27 4:32 4:38 4:43 4:49 4:54 4:6  4:65 4:70
4:11 4:17 4:22 4:28 4:33 4:39 4:44 4:5  4:55 4:60 4:66 4:71
4:12 4:18 4:23 4:29 4:34 4:4  4:45 4:50 4:56 4:61 4:67 4:8
4:13 4:19 4:24 4:3  4:35 4:40 4:46 4:51 4:57 4:62 4:68 4:9
```

Figure 18. Example listing of a Vmhba directory for a VMware ESX server

Note: In a multipathing configuration, only 1 vmhba directory displays entries for the shared device.

Each of the entries provides information on the drives that it points to, as shown in Figure 19:

```
[root@denethor root]# cat /proc/vmware/scsi/vmhba0//4\:71
Vendor: IBM      Model: 2105800    Rev: 1.89
Type:   Direct-Access          ANSI SCSI revision: 03
Size:   24795 Mbytes
Queue Depth: 16

Partition Info:
Block size: 512
Num Blocks: 50781312

      num:      Start      Size      Type
      1:         63 50765337      fb

Partition      VM      cmds      reads      KBread      writes      KBwritten      cmdsAbrt
busRst      paeCmds      paeCopies      splitCmds      splitCopies      issueAvg      totalAvg
0      -      168      87      46      0      0 0      0
0      0      0      0      0      20788      936149      0
1      -      8837      8581      33869      128      64 0      0
0      0      0      0      0      14038      762180      0

      VM      Shares      cmds      reads      KBread      writes      KBwritten      cmdsAbrt      busRst
paeCmds      paeCopies      splitCmds      splitCopies      issueAvg      totalAvg      active      queued
virtTime
127      1000      9005      8668      33915      128      64      0      0
0      0      0      0      0      14164      765426      0      0
8728000000
Total      1000      9005      8668      33915      128      64      0      0
0      0      0      0      0      14164      765426      0      0
8728000000

Paths:fixed
vmhba0:4:71 on*#
vmhba1:4:71 on

Active: 0 Queued: 0
```

Figure 19. Example of Vmhba entries for a VMware ESX server

Note: The last section displays information on multipathing; refer to VMware documentation for more information on the multipathing configuration. The storage unit can interact with both FIXED and MRU policies.

Disk device discovery on VMware ESX

This section provides information about disk device discovery on VMware ESX.

To rediscover LUNs from the Console OS use the command `vmkfstools -s vmhbaC` where C is the number of the virtual host adapter to rescan.

Persistent binding

Persistent binding is highly recommended in a SAN environment especially when several storage subsystems are available to the VMware ESX Server.

Persistent binding means that a device has the same identification to the operating system after it restarts and after other devices are added to the operating system.

In SAN, the storage SCSI ID is dictated by the order in which the VMware ESX server discovers the subsystems at system start time. In some instances, this order might differ because of a device addition or removal. This can cause VMware ESX server to assign a different target ID and change the device addressing (`vmhba <C> : <T> : <D>`). This can cause the virtual machines to report their storage as missing when the system is starting. Persistent binding forces the VMware ESX server to always use the same target ID for a storage unit no matter of the discovery order.

To enable persistent binding on all adapters, use the command: **`pbind.pl -A`**.

Note: You do not need to use this command if you are using v3.x.

Configuring an Intel or AMD host running VMware ESX Server

This section describes how you use the `fdisk` utility to partition a server disk and how you create a file system on an Intel or AMD host running VMware ESX Server.

Note: IBM does not support VMware ESX running on system LUN/disk or extended VMFS2 file system.

The following requirement is only necessary when the different host adapters are located on the same VMware ESX Server, not when devices are shared between VMware ESX Servers. VMware ESX multipathing requires that the storage unit present the disk drive modules in the same LUN sequence for each fibre-channel adapter on the host. When you assign volumes to target hosts using the DS Storage Manager, select the **use same ID/LUN in source and target** option.

Partitioning storage unit disks for an Intel or AMD host running VMware ESX Server

This section provides information about partitioning storage unit disks for an Intel or AMD host running VMware ESX server.

For ESX v3.x servers, use the VMware Virtual Center product for LUN partitioning and storage management.

Before you create a virtual machine file system (VMFS) file system, partition the disk by using the `fdisk` utility. You have to specify the special device file of the disk that you want to partition when you run the `fdisk` utility. The type of the partition that is used by virtual machine file system (VMFS) must be `fb` (must be set with the `fdisk` command).

To identify the special device file to use, run the command: `vmkfstools -N vmhba<C>:<T>:<D>:0`, where <C> is the virtual host adapter number, <T> is the SCSI ID of the storage unit, and <D> is the LUN number of the disk.

```
[root@denethor root]# vmkfstools -N vmhba0:4:60:0  
/dev/vsd60
```

Figure 20 provides an example of different fdisk utility options and Figure 21 on page 135 provides an example of primary partition on the disk /dev/vsd71.

```
[root@denethor root]# fdisk /dev/vsd71  
  
The number of cylinders for this disk is set to 3160.  
There is nothing wrong with that, but this is larger than 1024,  
and could in certain setups cause problems with:  
1) software that runs at boot time (e.g., old versions of LILO)  
2) booting and partitioning software from other OSs  
   (e.g., DOS FDISK, OS/2 FDISK)  
  
Command (m for help): m  
Command action  
  a toggle a bootable flag  
  b edit bsd disklabel  
  c toggle the dos compatibility flag  
  d delete a partition  
  l list known partition types  
  m print this menu  
  n add a new partition  
  o create a new empty DOS partition table  
  p print the partition table  
  q quit without saving changes  
  s create a new empty Sun disklabel  
  t change a partition's system id  
  u change display/entry units  
  v verify the partition table  
  w write table to disk and exit  
  x extra functionality (experts only)
```

Figure 20. Example of the different options for the fdisk utility on a VMware ESX server

```

Command (m for help): n
Command action
  e   extended
  p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-3160, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-3160, default 3160):
Using default value 3160

Command (m for help): p

Disk /dev/vsd71: 255 heads, 63 sectors, 3160 cylinders
Units = cylinders of 16065 * 512 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/vsd71p1            1         3160    25382668+   83   Linux

Command (m for help): t
Partition number (1-4): 1
Hex code (type L to list codes): fb
Changed system type of partition 1 to fb (Unknown)

Command (m for help): p

Disk /dev/vsd71: 255 heads, 63 sectors, 3160 cylinders
Units = cylinders of 16065 * 512 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/vsd71p1            1         3160    25382668+   fb  Unknown

```

Figure 21. Example of primary partition on the disk /dev/vsd71 for a VMware ESX server

Creating a file system for an Intel or AMD host running a VMware ESX server

This section provides information about creating and using VMFS file system for an Intel or AMD host running a VMware ESX server.

Note: If you are using v3.x, use *vmfs3* in the commands below instead of *vmfs2*.

After you partition the disk as described in “Partitioning storage unit disks for an Intel or AMD host running VMware ESX Server” on page 133, the next step requires that you create a VMFS file system by using the following commands:

```
vmkfstools -C vmfs2 <device>
```

<device> is on the format vmhba <C>:<T>:<D>:<P>, where <C> is the virtual host adapter, <T> is the SCSI ID of the storage unit, <D> is the Lun ID, and <P> is the partition number to use on that device (previously created with the fdisk utility). In most cases it will be 1.

```

[root@denethor vmfs]# vmkfstools -C vmfs2 vmhba0:4:71:1
This is already a file system on vmhba0:4:71:1. Overwrite it (y/n)? y
Creating file system on vmhba0:4:71:1 with blockSize 1048576

```

After this operation is completed, the new file system is under /vmfs.

```
vmkfstools -N <label name>
```

Note: This step is optional but we strongly recommended that you do it because it can ease your task of storage management. It can become a mandatory step if you do not configure persistent binding on the VMware ESX server.

Copy Services considerations for VMware ESX servers

This section provides considerations for using the Copy Services functions with VMware ESX servers.

Considerations for ESX v2.x

IBM supports FlashCopy and Metro Mirror functions with the following limitations:

- Copy Services functions are only available for individual LUNs on the DS6000.
- The source and target volumes must exist on separate physical machines.
- Remote Mirror and Copy read from secondary option is not supported.
- Rescan of LUNs might be required.

Copy Services are only available for individual LUNs.

The DS6000 does not have any knowledge of VMFS volumes or extents so you must copy all LUNs that might be contained within a VMFS volume. To prevent LUN locking problems, it is mandatory that no virtual machine uses a copy target volume during the copy process.

The source and target volumes must be on separate physical machines

The VMFS file system uses a label mechanism and if it encounters two volumes with the same label, only one remains accessible.

Remote Mirror and Copy read from secondary option is not supported

“Read from secondary” allows a target volume to be opened in read-only mode during a Remote Mirror and Copy operation. VMware ESX Server does not open volumes in read-only mode and therefore is not able to access a Remote Mirror and Copy target while a copy is in progress.

Rescan of LUNS might be required

VMFS is able to discover changes on a hard drive at the file system level. It can manage a copy as long as the system knows that there is a VMFS file system to discover:

- If the copy target volume was discovered at system start time without a VMFS file system, you must perform a rescan operation for the changes.
- If there is a VMFS file system, the server can use the new information immediately without any intervention.

Considerations for ESX v3.x

Copy Services functions are supported for any DS6000 LUN that is attached to an ESX server, with the following important restrictions and recommendations:

- The copy relationship must not be persistent if the target LUNs are used with a VMware host. If you are using the Metro Mirror function, the complete relationship must be severed; if the suspend or read-only options are used, VMware cannot use the copied LUNs.
- If the target LUNs are used within the same VMware Datacenter as the source LUNs, a resignature operation is required (see the VMware SAN Configuration Manual, version 3.x, that is available at <http://www.vmware.com>).

- Ensure that you create separate DS6000 volume groups for the set of source LUNs and the sets of target LUNs. The LUNs in the target volume group must have the same order and size as LUNs in the source volume group. Only one set of target LUNs can be presented to a single VMware host.
- Ensure that LUN names are distinctive for each type or set of copy operations. Distinctive names aid in postoperation identification of LUNs.
- Raw device mapping (RDM) LUNs might require reattachment to the OS image that has been copied. The OS image must be edited so that the existing RDM LUNs are removed, and then add the appropriate LUNs. Data is preserved during this operation.
- If an OS image is running during a copy operation, there might be a loss of data due to the timing of I/O buffer flushes. For absolute data integrity, shut down the guest OS image during the copy operation.

Intel or AMD running Microsoft Windows 2000 or 2003 host attachment

This chapter describes how you attach an Intel or Advanced Micro Devices (AMD) host system running Windows 2000 or 2003 to a DS6000 with fibre-channel adapters.

This chapter contains the following host attachment sections:

- Installing and configuring the Emulex adapter
- Installing and configuring the Netfinity adapter
- Installing and configuring the QLogic adapter
- Updating the Windows 2000 or 2003 device driver
- Locating the WWPN for a Windows 2000 or 2003 host
- Verifying that Windows 2000 or 2003 is configured for storage
- Setting the TimeoutValue registry
- Installing remote fibre-channel boot support

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Version 1.4.0.0 or later of the SDD supports Windows Server 2000 or 2003 host system in a clustering environment. To have failover protection on an open system, the SDD requires at least two paths, and SDD allows a maximum of 32 paths. See <http://www.ibm.com/systems/support/storage/software/sdd> for the most current information about SDD.

Installing an Emulex adapter and driver for Windows 2000 or 2003

This section describes how you install and configure an Emulex adapter card and driver for a Windows 2000 or 2003 host.

Single- and dual-port, fibre-channel interfaces with an Emulex adapter card support the following public and private loop modes:

- Target
- Public initiator
- Private initiator
- Target and public initiator
- Target and private initiator

Note: The Emulex LP9402DC adapter card is a 3.3 volt card only and requires a 3.3 volt PCI slot

Complete the following steps to install and configure an Emulex adapter:

1. Record the IEEE number that is printed on the adapter card. You can use the IEEE number to determine the WWPN.
2. Install and configure the adapter and driver on the host. See the installation instructions that are provided by your host adapter vendor for specific adapter and adapter driver instructions. Go to “Downloading and installing a host adapter driver” on page 30 or <http://www.emulex.com/ts/docoem/ibm/index.htm>.

Table 8 provides the recommended settings for the following Emulex adapters: LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, LP10000DC, LP11000, LP11002, LPE11000, and LPE11002 adapters. Parameters are provided for the Exlcfg and HBAnywhere interfaces.

3. Connect the cable to the adapter and to the DS6000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS6000, see Host configuration using the DS Storage Manager.

Table 8. Recommended configuration file parameters for the Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, LP10000DC, LP11000, LP11002, LPE11000, and LPE11002 adapters

Parameters - Exlcfg	Parameters - HBAnywhere	Recommended settings
Automatically map SCSI devices	Automap	Checked (enabled)
Query name server for all N-ports	NsQueryAll	Checked (enabled)
Allow multiple paths to SCSI targets	MultipleScsiClaims	Checked (enabled)
Register for state change	RegisterForStateChange	Checked (enabled)
Use report LUNs	UseReportLuns	Checked (enabled)
Use name server after RSCN	BrokenRscn	Checked (enabled) Obsolete for HBAnywhere (see note 1 on page 140)
BusResetForBreakRes	-	(see note 2 on page 140)
LUN mapping	MapLuns	Checked (enabled)
Automatic LUN mapping	MapLuns	Checked (enabled)
Scan in device ID order	ScanDeviceIdOrder	Not checked (disabled)
Disable target reset for tape devices	TapeResetMask	Not checked (disabled)
Use PLOGI instead of PDISC after LIP	ForcePlogi	Checked (enabled)
Enable class 2 for SCSI devices	Class2Enable	Not checked (disabled)

Table 8. Recommended configuration file parameters for the Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, LP10000DC, LP11000, LP11002, LPE11000, and LPE11002 adapters (continued)

Parameters - Exlcfg	Parameters - HBAnywhere	Recommended settings
Report unknown SCSI devices	ShowUnknowns	Not checked (disabled)
Look for disappearing devices	AttemptReEstablish	Not checked (disabled)
Translate queue full to busy	TranslateQueueFull	Checked (enabled) if attached LUNs are part of a PPRC or FlashCopy pair. Not checked (disabled) if attached LUNs are not part of a PPRC or FlashCopy pair.
Use bus reset status for retries	RegisterForStateChange	Not checked (disabled)
Retry unit attention	RetryUnitAttention	Not checked (disabled)
Retry PLOGI open failures	RetryPlogiOpenFailure	Not checked (disabled)
Enable FCP-2 Recovery	FcTape	Checked (enabled)
Enable FDMI	EnableFDMI	Not checked (disabled)
Maximum number of LUNs	MaximumLun	Equal to or greater than the number of the storage unit LUNs available to the host bus adapter
Maximum queue depth	MaximumQueueDepth	8
Driver Trace Mask 0x	ffTraceMask	0
Timer Settings (Tuning Drop Down Menu)		
Link Timer	LinkDownTimer	30 seconds (see note 3 on page 140)
Retries	RetryAttempts	64 (see note 3 on page 140)
E_D_TOV	EDTOV	2000 milliseconds (see note 3 on page 140)
AL_TOV	EDTOV	15 milliseconds (see note 3 on page 140)
Wait ready timer	WaitReadyTimerVal	45 seconds (see note 3 on page 140)
Retry timer	RetryTimer	2000 milliseconds (see note 3 on page 140)
R_A_TOV	RATOV	2 seconds (see note 3 on page 140)
ARB_TOV	ARB_TOV	1000 milliseconds (see note 3 on page 140)
Tuning Link Control		

Table 8. Recommended configuration file parameters for the Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, LP10000DC, LP11000, LP11002, LPE11000, and LPE11002 adapters (continued)

Parameters - Exlcfg	Parameters - HBAnywhere	Recommended settings
Topology	InitLinkFlags	<ul style="list-style-type: none"> • Point-to-point No 2 Gb/s 0x02000002 • Point-to-Point Yes 2 Gb/s 0x02800002 • Point-to-Point No 4 Gb/s 0x04000002 • Point-to-Point Yes 4 Gb/s 0x04800002 • Arbitrated Loop No 2 Gb/s 0x02000000 • Arbitrated Loop Yes 2 Gb/s 0x02800000 • Arbitrated Loop No 4 Gb/s 0x04000000 • Arbitrated Loop Yes 4 Gb/s 0x04800000 • No/Yes refers to autodetection of fabric • (see note 4)
Link speed	InitLinkFlags	See above
Host pointers in SLIM	Unable to find corresponding value	Not checked (disabled)
<p>Notes:</p> <ol style="list-style-type: none"> 1. These fields are no longer used, but have been retained for possible future use. 2. This parameter cannot be managed directly through the elxcfg interface. The format of this parameter is documented in the readme.txt file and in the following text: If 1, it will issue a bus reset instead of a target reset when receiving an IOCTL_STORAGE_BREAK_RESERVATION IOCTL. The default is 0, which will issue a target reset to a given target ID. 3. All timer-related parameters in the HBAnywhere interface are managed in the elxcfg interface under the Tuning/Timers menu. 4. Topology and link speed are managed in the elxcfg interface under the Tuning/Link control menu. In the HBAnywhere interface, they are combined into a single parameter: InitLinkFlags. The format of this parameter is documented in the readme.txt file and in the following text: <ul style="list-style-type: none"> • The InitLinkFlags is divided into several fields. Starting from bit 0, bits 0-2 are the initialization type. If bits 0-2 are 000, the link will be initialized as Arbitrated Loop. If bits 0-2 are 010, the link will be initialized as Point to Point. The presence of a fabric is automatically detected in either case. • Bit 23, if set, indicates the adapter should attempt the other topology type should the first selection fail to initialize. This option is only available with firmware that supports this option. • Bits 24 to 31 are the encoded speed value. If bits 24-31 are 0x00, the link speed is discovered automatically. If bits 24-31 are 0x01, the link speed is forced to 1 Gigabit per second. If bits 24-31 are 0x02, the link speed is forced to 2 Gigabits per second. If bits 24-31 are 0x04, the link speed is forced to 4 Gigabits per second. This option is only available with adapters that support more than one link speed. 		

Use Table 9 as a guideline for the recommended settings for the StorPort Miniport driver.

Table 9. StorPort Miniport driver settings

Parameter	Recommended Settings
AutoMap	2 (WWPN)
Class	3 (Class 3 FC)
Coalesce MsCnt	0
CoalesceRspCnt	8
DiscoveryDelay	0
EnableAck0	Enabled
EnableAck0	0
FrameSizeMSB	0
HardAPIa	0x0
InitTimeOut	15
LinkSpeed	4Gb or 2Gb
LinkTimeOut	30
LogErrors	3
NetworkOption	0
NodeTimeOut	30
NumFcpContext	512
QueueDepth	32
QueueTarget	0
RmaDepth	16
ScanDown	Enabled
TargetOption	0
Topology	<ul style="list-style-type: none"> • 0 Loop • 1 F_Port Fabric • 2 Loop then F_Port • 3 F_Port then Loop
TraceBufSiz	250000
UseAdisc	Disabled

Installing a Netfinity adapter and driver for Windows 2000 or 2003

This section provides instructions for installing and configuring the Netfinity[®] adapter card and driver for the Windows 2000 or 2003 host system.

1. Obtain an optical multimode cable with a standard-connector style of duplex connector.
2. Write down the serial number of the adapter card.
In the unlikely event that the nonvolatile random access memory is damaged, the system will prompt you for the Netfinityadapter serial number.
3. Check the system board and make any configuration changes necessary to accommodate the Netfinity adapter.

4. Turn off the peripheral devices and then turn off the host.
5. Disconnect the power cord of the host.
6. Install the card in the host.
 - a. Remove the cover to the host.
 - b. Locate an available 64-bit PCI slot. If a 64-bit slot is not available, use a 32-bit PCI slot.
 - c. Align the adapter with the expansion slot, and carefully press the adapter into the slot until you hear the adapter snap into place.
 - d. Secure the adapter.
 - e. Connect one end of the fiber cable to the J1 connector on the adapter and connect the other end to the DS6000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS6000, see Host configuration using the DS Storage Manager.
 - f. Carefully reinstall and secure the cover.
 - g. Connect the host power cables.
7. Turn on all external fibre-channel devices.
8. The host monitor displays start up information. Watch this information for messages about the adapter.

Note: You will not see the BIOS menu because the adapter BIOS is set to disabled by default. .

9. To see the devices, press Alt+Q, and then use the utility program.
10. Install and configure the adapter driver on the host. See the installation instructions that are provided by your host adapter vendor for specific adapter and adapter driver instructions. Go to "Downloading and installing a host adapter driver" on page 30.

Note: Retrieve a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>.

Documentation that comes with your Netfinity adapter

This section lists documentation that comes with your Netfinity adapter card.

The following files are on the diskette that is shipped with your adapter:

Type	Description
TXTSETUP.OEM	The driver installation script for the initial Windows text setup
QLOGIC	The identification file for the Windows setup program.
\W2K\OEMSETUP.IN	The driver installation script for the Windows setup program
\W2K\QL2200.SYS	The Windows 2000 device driver
\W2K\README.TXT	Helpful hints about the Windows 2000 device driver
\W2K\RELEASE.TXT	The history of release fixes

Installing the Netfinity device driver for Windows 2000 or 2003

This section provides instructions to install the Netfinity device driver on a host that already has Windows 2000 or 2003 installed.

You must already have installed the Windows 2000 operating system installed on your host machine.

1. Install the Netfinity adapter.
2. Restart Windows.
3. Insert the Netfinity host adapter device driver for Windows 2000 or 2003 diskette and follow the instructions that are displayed.

Installing a QLogic adapter and driver for Windows 2000 or Windows 2003

This section describes how to install a QLogic adapter card on a Windows 2000 or Windows 2003 host and provides the recommending settings for the adapter.

1. Install the QLogic adapter card in the host system.
2. Connect the cable to the storage unit port.
3. Install and configure the adapter driver. See the installation instructions that are provided by your host adapter vendor for specific adapter and adapter driver instructions. Go to “Downloading and installing a host adapter driver” on page 30 or to http://support.qlogic.com/support/oem_ibm.asp.
4. Restart the host.
5. Press Ctrl+Q to get to the **FAST!Util** menu.
6. Use FAST!Util to set the parameters to the values that are listed in Table 10.

Note: “Installing a QLogic adapter and driver for Windows 2000 or Windows 2003” contains settings for the QLogic QLA23xx, QLA246x, and QLE246x adapters.

Table 10. Recommended settings for the QLogic QLA23xx, QLA246x, and QLE246x adapters for Windows 2000 or Windows 2003

Parameters	Recommended settings
Basic Input/Output System (BIOS)	Disabled
Frame size	2048
Loop reset delay	5 (minimum)
Adapter hard loop ID	Disabled
Hard loop ID	0
Execution throttle	100
Fast command posting	Enabled
>4 GB addressing	Disabled (for 32-bit systems)
LUNs per target	0
Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes Note: Enable LIP reset, Enable LIP full login, and Enable target reset parameters control the behavior of the adapter when Windows attempts a SCSI bus reset. You must perform a target reset operation to make cluster failovers work. Use the SCSI bus device reset option to clear SCSI reservations.
Login retry count	20 (minimum)
Port down retry count	20 (minimum)

Table 10. Recommended settings for the QLogic QLA23xx, QLA246x, and QLE246x adapters for Windows 2000 or Windows 2003 (continued)

Parameters	Recommended settings
Driver load RISC code	Enabled
Enable database updates	No
Disable database load	No
IOCB allocation	256
Extended error logging	Disabled (You might want to enable this parameter if you use the debugging function.)
Type of connection	<ul style="list-style-type: none"> • Loop only • Point-to-point • Loop preferred (If you cannot use arbitrated loop, default to point-to-point.) • Point-to point; otherwise loop only (If you cannot use point-to-point, default to arbitrated loop.) <p>Notes:</p> <ol style="list-style-type: none"> 1. If you connect the storage unit directly to the host system, the option that you select must match the port connections on the storage unit. 2. If you connect through a switch, the options do not need to match the port connections because the storage unit is point-to-point. 3. The appropriate host adapter on the server must also support a point-to-point connection on a direct connection. 4. If you use adapter cards from different manufacturers, they will not function properly in a direct point-to-point connection. This is not true if you connect through a switch, because the storage unit is point-to-point.

7. Restart the server.

Note: If you are using Windows 2003, QLA2340/2342, and an IBM 2005-B32 (Brocade Silkworm 4100), the switch might report false errors due to a timing mismatch between the HBA and the switch during server reboot. This makes diagnosis difficult. To avoid this, move the default jumper positions on the HBA card. Normal settings on the QLA2340/2340L HBAs are described in Section 2.1.2 of the *Qlogic FC5054601 User's Guide for the Qlogic QLA2340/2342*. See the location of the jumpers in Figure 22 on page 145, and move the jumpers as indicated in Table 11 on page 145. This shuts off the HBA laser during reboot, prevents port faults, and allows the server to detect SAN storage.




Figure 22. Jumper location on the QLA2340

Table 11. Jumper change for QLA2340/2342, and an IBM 2005-B32 (Brocade Silkstorm 4100) running Windows 2003

	Jumper plug	Jumper
Factory Default	pins 2-3	J3 jumper
Change	pins 1-2	J3 jumper

Updating the Windows 2000 or 2003 device driver

This section provides instructions to replace an existing Windows 2000 or 2003 device driver with an updated version.

1. Right click **My Computer** → **Manage** → **Device Manager**.
2. Click the plus (+) sign to expand the menu for **SCSI and RAID Controllers**.
3. Right click on the adapter that you want to update and select **Properties**.
4. Select the tab for **Driver**.
5. Click **Update Driver** and follow the instructions that appear on the screen.

Locating the WWPN for a Windows 2000 or 2003 host

Complete this task to locate the worldwide port name (WWPN) for a Windows 2000 or Windows 2003 with either the QLogic or the Emulex adapter.

1. Perform the following steps to locate the WWPN for a Windows 2000 or 2003 host system with a QLogic adapter:
 - a. Restart the server.
 - b. Press Alt+Q to get the **FAST!Util** menu.

Note: If you have more than one fibre-channel adapter installed, all the fibre-channel adapters display. Scroll down to the adapter you want. Press Enter.

- c. From the **FAST!Util** menu, scroll down and select **Select Host Adapter**.
- d. Scroll up and highlight **Configuration Settings**. Press Enter.
- e. From the **Configuration Settings** menu, click **Host Adapter Settings**.
- f. Write down the 16-digit alphanumeric string that is displayed on your screen.

2. Perform the following steps to locate the WWPN for a Windows 2000 or 2003 host system with an Emulex adapter:
 - a. Click **Start** → **Programs** → **Emulex Configuration Tool**.
 - b. From the Emulex Configuration Tool menu in the Available Adapters window, double-click the adapter entry for which you want to display the WWPN information.

Verifying that Windows 2000 or 2003 is configured for storage

This section provides instructions to determine if your Windows 2000 or 2003 host system is configured for storage.

1. Partition new drives with Disk Management.
2. From the Windows Server desktop, right click **Start**.
3. Click **Explore** and verify that you can see the fibre-channel drives.
4. Select a large file (for example, 9 MB file), and drag (copy) it to a fibre-channel drive.
5. Verify the file copied to the selected drive.

Setting the TimeOutValue registry for Windows 2000 or 2003

This section describes how to ensure optimum availability and recoverability when you attach a storage unit to a Windows 2000 or 2003 host system by setting the Windows TimeOutValue registry.

You must set the TimeOutValue value associated with the supported host adapters to 60 seconds. The setting is consistent with the configuration for FC-AL adapters and disk subsystems when attached to Windows 2000 or 2003 host system.

The host adapter uses the TimeOutValue parameter to bind its recovery actions and responses to the disk subsystem. The value exists in different places in the system configuration. You can retrieve and use it in different ways depending on the type of host adapter that is installed.

Follow these steps to modify the value in the Windows 2000 or 2003 registry.

1. From the **Run** menu or command prompt, type:
Regedt32.exe
2. Navigate to the following registry key:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk
3. Look for the value called TimeOutValue. If the value called TimeOutValue does not exist, go to step 3a. If the TimeOutValue exists, go to step 4.
 - a. Click **Edit** → **Add Value...**
 - b. For ValueName, type: TimeOutValue.
 - c. For data type, type: REG-DWORD.
 - d. Click **OK**.
 - e. For data, type: 3c.
 - f. For radix, click **Hex**.
 - g. Click **OK**.
4. If the value exists and is less than 0x0000003c (60 decimal), perform the following steps to increase it to 0x3c.
 - a. Click **TimeOutValue**.
 - b. Click **Edit** → **DWORD...**

- c. For data, type: 3c.
- d. For radix, click **Hex**.
- e. Click **OK**.
5. Exit the Regedt32 program.
6. Restart your Windows 2000 or 2003 server for the changes to take effect.

Installing remote fibre-channel boot support for a Windows 2000 or 2003 host system

This section provides instructions for installing remote fibre-channel boot support for a Windows 2000 or 2003 host system.

You must obtain the following items before you can begin to install remote fibre-channel boot support:

- The *IBM System Storage Enterprise Storage Server Subsystem Device Driver User's Guide* at <http://www.ibm.com/systems/support/storage/software/sdd>.
- The *IBM System Storage DS6000 Introduction and Planning Guide*.
- Either the Windows 2000/2003 CD or a set of Windows 2000/2003 setup floppy disks with the Windows 2000/2003 CD.
- The Windows 2000 or 2003 installation instructions.
- The supported fibre-channel host adapter driver package.
- The supported fibre-channel host adapter driver documentation.

Ensure that the host adapters are installed and cable connections of the host, switch, and storage unit are complete and functional.

Extract and copy the driver files to a floppy.

Remote fibre-channel boot support is available for Windows 2000 or 2003 with or without SDD. For information on the SDD hardware and software installation requirements, see *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Configure zoning and obtain storage

1. Reboot the host system to obtain the Host Bus Adapter WWPN:
 - a. **For QLogic host adapters:** A QLogic BIOS banner appears, press <Alt Q> or press <Ctrl Q> as indicated, depending on the BIOS level. For each host adapter, select Configuration Settings -> Host Bus Adapter Settings and write down the host adapter WWPN.
 - b. **For Emulex host adapters:** When the BIOS banner appears, press either <Alt E> or press <Alt F5> as indicated, depending on the BIOS level. A list of adapters is displayed. For each host adapter, select the adapter and write down the host adapter WWPN.
2. Configure appropriate zoning on the switch to include only the storage unit WWPN and one host adapter WWPN.
3. Use the IBM System Storage DS Storage Manager to add the host system to the list of host systems that are available for attachment to the storage unit. For detailed information of how to use the IBM System Storage DS Storage Manager, refer to the *IBM System Storage DS6000 Introduction and Planning Guide*.
4. Use the IBM System Storage DS Storage Manager to perform storage allocation to the host system.

Flash QLogic host adapter

Because the host adapter BIOS is responsible for scanning the fibre-channel bus to locate the remote boot device, check that the latest supported host adapter BIOS is loaded. If not, follow the host adapter documentation to update the BIOS of the host adapter.

Configure QLogic and Emulex host adapters

1. Return to the host system.
2. **For QLogic host adapters:** When the QLogic BIOS banner appears, press <Ctrl Q>. The QLogic Fast!UTIL Menu appears. Depending on the host adapter BIOS version, the BIOS menu may be different in the following step:
 - a. Select Configuration Settings->Host Bus Adapter Settings. Set only one of the Host Adapters BIOS to **Enabled** and set **Connection Options** appropriately. For example, set to 1 for fabric point-to-point. Press ESC to return to the previous menu.
 - b. Select Configuration Settings->Selectable Boot Settings. Set Selectable Boot to **Enabled** and select the appropriate Boot Port Name and the LUN that corresponds to the WWPN of the storage unit.
3. **For Emulex host adapters:** When Emulex BIOS banner appears, press Alt E or press Alt F5 as indicated, depending on the BIOS level. A list of adapters is displayed. For each host adapter, select the adapter by entering the appropriate number.
 - a. Set only one host adapter to **Enabled** and set **Topology** appropriately. For example, set to 1 for fabric point-to-point. Press ESC to return to the previous menu.
 - b. Select Configure Boot Devices and configure one boot device, including the boot device, the boot entry WWPN, and the starting LUN.
4. Exit the BIOS utility and select to save changes.

Windows 2000 and 2003 installation

Apply the Windows Server installation instructions. These instructions include the following steps:

1. Reboot the host system and insert the Windows setup CD in the CD drive.
2. If prompted with the message, **Press any key to boot from CD**, press any key.
3. Respond to **Press F6 if you want to install a third party SCSI or RAID driver** message by pressing F6
4. On **Welcome to Setup** panel, press Enter to select the **Set up of Windows now** message.
5. Answer appropriately on the **Windows Licensing Agreement** panel.
6. Press ESC to reply to the **Continue installing a fresh copy of Windows** prompt.
7. When prompted, choose the disk on the storage unit where the Windows Server operating system will reside.

Windows 2000 and 2003 Post Installation

1. Once Windows Server installation is complete, install the IBM Subsystem Device Driver (SDD). Refer to the *IBM System Storage Multipath Subsystem Device Driver User's Guide* and the SDD Readme.txt for the installation of SDD.
2. Update zoning on the switch to permit the availability of multiple paths.

Intel running Novell NetWare host attachment

This chapter describes how you attach a Novell NetWare host system to a DS6000 with fibre-channel adapters.

This chapter contains the following host attachment topics:

- Downloading and installing an Emulex or QLogic adapter
 - Installing the Emulex adapter for a Novell NetWare host
 - Downloading the current Emulex adapter driver for a Novell NetWare host
 - Installing the Emulex adapter driver for a Novell NetWare host
 - Downloading the current QLogic adapter driver for a Novell NetWare host
 - Installing the QLogic QLA23xx adapter for a Novell NetWare host
 - Installing QLogic adapter drivers for a Novell NetWare host
- Locating the WWPN for a Novell NetWare host

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Installing an Emulex adapter and driver for a Novell NetWare host

This section provides installation steps for installing the Emulex adapter driver on a Novell NetWare host.

If you are installing the fibre-channel adapter for the first time, you must specify the correct topology. You must also select the appropriate device mapping driver.

Perform the following steps to install an Emulex adapter and driver:

1. Record the IEEE number that is printed on the adapter card. You can use the IEEE number to determine the WWPN.
2. Install the Emulex adapter card in the host system.
3. Connect the cable to the adapter and to the DS6000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS6000, see Host configuration using the DS Storage Manager.
4. Install and configure the adapter driver on the host. Refer to the installation instructions provided by your host adapter vendor for specific adapter and adapter driver instructions. Go to “Downloading and installing a host adapter driver” on page 30 or <http://www.emulex.com/ts/docoem/ibm/index.htm>.
5. Once you have installed the driver from the NetWare server console, enter the following command: `nwconfig`. This command applies to NetWare 5.0 or later.
6. Click **NCF Files Options**.
7. Select **Edit STARTUP.NCF** file.
8. Insert the following statement at the beginning of the STARTUP.NCF file: `SET MULTI-PATH SUPPORT=OFF`

9. Find the LOAD statement for the LPFC.HAM driver and ensure that it reads as follows: `LOAD LPFC.HAM SLOT=XXXX TOPOLOGY=1 /LUN MAX_LUN=32 RSCNZONE=1`
10. Save the `STARTUP.NCF` file.
11. Restart the server, or unload and reload the LPFC driver manually.

Note: The `MAX_LUN` parameter must be set to a number greater than the LUNs you have assigned to the adapter.

Installing a QLogic adapter and driver for a Novell NetWare host

This section describes how you install a QLogic adapter and driver on a Novell NetWare host system.

Single- and dual-port fibre-channel interfaces with the QLogic QLA23xx and QLA246x adapters support the following public and private loop modes:

- Target
- Public initiator
- Private initiator
- Target and public initiator
- Target and private initiator

Perform the following steps to install a QLogic adapter and driver:

Note: The following steps are an example of a configuration using the QLogic QLA23xx adapter. The configuration for your adapter might differ.

1. Install the QLogic adapter in the host system.
2. Connect the cable to the adapter and to the DS6000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS6000, see Host configuration using the DS Storage Manager.
3. Restart the server.
4. Press `Ctrl+Q` to get to the **FAST!Util** menu.
5. Use **FAST!Util** to set the parameters to the values listed in Table 12:

Table 12. Recommended settings for the QLogic QLA23xx, and QLA246x adapters for a Novell NetWare host

Parameters	Recommended settings
Basic Input/Output System (BIOS)	Disabled
Frame size	2048
Loop reset delay	5 (minimum)
Adapter hard loop ID	Disabled
Hard loop ID	0
Execution throttle	16
Fast command posting	Enabled
>4 GB addressing	Disabled (for 32-bit systems)
LUNs per target	32 Note: Set LUNs per target to a number greater than the LUNs you have assigned to the adapter.
Enable LIP reset	No

Table 12. Recommended settings for the QLogic QLA23xx, and QLA246x adapters for a Novell NetWare host (continued)

Parameters	Recommended settings
Enable LIP full login	Yes
Enable target reset	Yes
Login retry count	20 (minimum)
Port down retry count	20 (minimum)
Driver load RISC code	Enabled
Enable database updates	No
Disable database load	No
IOCB allocation	256
Extended error logging	Disabled (You might want to enable this parameter if you use debugging.)
Type of connection	<ul style="list-style-type: none"> • Loop only • Point-to-point • Loop preferred (If you cannot use arbitrated loop, then default to point-to-point.) • Point-to point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop.) <p>Note: If you connect the storage unit directly to the host system, the option you select must match the port connections on the storage unit.</p>

6. Install and configure the adapter driver on the host. See “Installing a QLogic adapter driver for a Novell NetWare host” for instructions.
7. Restart the server.

Installing a QLogic adapter driver for a Novell NetWare host

This section provides instructions for installing a QLogic adapter driver for a Novell NetWare host.

1. From the NetWare server console, type `nwconfig` for NetWare 5.0 or later.
2. Click **Driver Options**.
3. Click **Configure Disk and Storage Device Drivers**.
4. Click **Select an additional driver**.
5. Press the Insert key.
6. Insert a floppy diskette with the QLogic drivers into the A:\ drive of the NetWare server. Press Enter.
The available driver is displayed.

Note: See “Downloading and installing a host adapter driver” on page 30 for instructions on downloading a QLogic adapter driver.

7. Click the driver for the QLogic card, and press Enter.
8. Select **Select/Modify driver parameters**, and type the slot number of the QLogic card into the slot number parameter.
9. Set the Scan All LUNs parameter to Yes.
10. Press the Tab key and check the **Save Parameters and Load Driver** check box.
11. Press Esc twice to return to the **Configuration Options** menu.
 - a. Highlight the **NCF Files Options** menu and press Enter.

- b. Select **Edit the STARTUP.NCF** and press Enter.
- c. Verify the location of the STARTUP.NCF file, and modify it as necessary. Press Enter.
- d. Enter the following statement at the beginning of the STARTUP.NCF file:
SET MULTI-PATH SUPPORT=OFF

- e. Find the LOAD statement for the QL2300.HAM driver and ensure that it reads as follows:
LOAD QL2300.HAM slot=xxxx /luns /xretry=12 /xtimeout=120

Note: If you are using the IBM Multipath Subsystem Device Driver (SDD), you must turn off QLogic failover at the adapter and portnames level using the following command: LOAD QL2300.HAM slot=xxxx /luns /xretry=12 /xtimeout=120 /allpaths /portnames.

- f. Find the line that reads:

LOAD QLCSFTE

Modify the line to read: REM LOAD QLCSFTE

- g. Press Esc and select **Yes** when prompted to save the file.

12. Type UNLOAD QL2300.HAM.

13. Type LOAD QL2300.HAM slot=xxxx /luns /xretry=12 /xtimeout=120, where xxxx is the PCI slot number where the QL23xx card is installed.

Note: If you are using the IBM Multipath Subsystem Device Driver (SDD), you must turn off QLogic failover at the adapter and portnames level using the following command: LOAD QL2300.HAM slot=xxxx /luns /xretry=12 /xtimeout=120 /allpaths /portnames.

The storage unit volumes are displayed in the devices list. Create volumes using the nssmu utility if necessary.

Locating the WWPN for a QLogic adapter on a Novell NetWare host

Complete this task to locate the WWPN for a Novell NetWare host.

1. Restart the server.
2. Press Alt+Q to get the **FAST!Util** menu.

Note: If you have more than one fibre channel adapter installed, all the adapters display on the screen. Scroll down to the adapter you want. Press Enter.

3. From the **FAST!Util** menu, scroll down and select **Select Host Adapter**.
4. Scroll up and highlight **Configuration Settings**. Press Enter.
5. From the **Configuration Settings** menu, click **Host Adapter Settings**.
6. Write down the 16-digit alphanumeric string that is displayed on your screen.

Silicon Graphics host attachment

This chapter describes how you attach a Silicon Graphics (SGI) host system to a DS6000 with fibre-channel adapters.

This chapter contains the following sections:

- Checking the version of the IRIX operating system
- Installing a fibre-channel adapter card for the SGI host system

- Installing an optical cable for SGI in a switched-fabric topology
- Installing an optical cable for SGI in an arbitrated loop topology
- Confirming switch connectivity for SGI
- Displaying zoning information for the switch
- Locating the WWPN for an SGI host
- Confirming storage connectivity on an SGI host
- Configuring the storage unit for host failover using IRIX multipathing
- Confirming storage in a switched-fabric or arbitrated loop topology

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Notes:

1. Ensure that you have the IRIX64 6.5.14 version or later installed on your host system.
2. The SGI host system does not support the SDD because IRIX manages multipathing.
3. You can also connect a nonclustered SGI host system to the storage unit using the switched fabric topology on optical fiber media, SGI fibre-channel host bus adapters, and the IBM 2109 switch. You can also connect directly from an SGI host system to a storage unit with the arbitrated loop topology.

Checking the version of the IRIX operating system

This section provides the command to check the IRIX operating system version.

To confirm that you have the correct version of the IRIX operating system installed on your host system, type `# uname -Rs`.

Installing a fibre-channel adapter card for the SGI host system

This section lists the steps to install a fibre-channel adapter card for the SGI host system.

1. Install the adapter card in the host system.
2. Connect the cable to the storage unit port.
3. Restart the server.

Note: The IRIX driver is part of the IRIX operating system. There are no additional driver packages. SGI is responsible for support and distribution of the adapter when you use it in an SGI host system.

Verifying the installation of a fibre-channel adapter card for SGI

Complete this task to verify your installation of a fibre-channel adapter card for SGI.

1. Type `# ls -l /hw/scsi_ctlr` to verify that the command created the links for the controllers.

The following is example output. In this example, the PCI host bus adapters are installed in slots 6 and 8, and the XIO host bus adapters are installed in slots 7 and 9.

```
# ls -l /hw/scsi_ctlr
lrw----- 1 root sys 50 Mar 27 10:31 6 -> /hw/module/1/slot/io7/xtalk_pci/pci/1/scsi_ctlr/0
lrw----- 1 root sys 48 Mar 27 10:31 7 -> /hw/module/1/slot/io2/pci_xio/pci/0/scsi_ctlr/0
lrw----- 1 root sys 50 Mar 27 10:31 8 -> /hw/module/1/slot/io6/xtalk_pci/pci/1/scsi_ctlr/0
lrw----- 1 root sys 48 Mar 27 10:31 9 -> /hw/module/1/slot/io2/pci_xio/pci/1/scsi_ctlr/0
```

2. To confirm the installation of the fibre-channel driver, type `scsiha -w /hw/scsi_ctlr/<N>/bus`, where `<N>` is the number of the PCI bus.

Configuring the fibre-channel adapter drivers for SGI

Complete this task to configure fibre-channel adapter drivers for SGI.

1. Type `# cd /var/sysgen/master.d` to navigate to the necessary directory.
2. Type `# vi qlfc` to open the file.
3. Edit and configure the adapter driver for switched-fabric or arbitrated loop topology. Use Table 13 as a guide to set the `qlfc_use_connection_mode` parameters.

Table 13. SGI `qlfc_use_connection_mode` parameters

<code>qlfc_use_connection_mode</code>	Topology
0	Arbitrated loop
1	Switched-fabric
2	Loop preferred (default)
3	Point-to-point preferred, then loop

Installing an optical cable for SGI in a switched-fabric topology

Use the commands in this section to install an optical cable for SGI in a switched-fabric topology.

1. Connect the fibre-channel host adapters from the SGI host system to an IBM 2109 or a Brocade switch.
2. Connect the storage unit fibre-channel host adapters to the switch.

Installing an optical cable for SGI in an arbitrated loop topology

Use the commands in this section to install an optical cable for SGI in an arbitrated loop topology.

1. Connect the fibre-channel host adapters from the SGI host system directly to the fibre-channel host adapters for the storage unit.
2. Connect the storage unit fibre-channel host adapters to the switch.

Confirming switch connectivity for SGI

Use the following commands to confirm switch connectivity for SGI.

1. Log in to the IBM or Brocade switch as administrator.
2. At the console prompt, type `switchshow`.

The following is example output. In this example, ports 9 and 11 are used for testing other devices.

```
snj2109s16h1:os1> switchshow
switchName:      snj2109s16h1
switchType:      2.4
switchState:     Online
switchRole:      Principal
switchDomain:     1
switchId:        fffc01
switchWwn:       10:00:00:60:69:12:30:25
switchBeacon:    OFF
port 0: sw Online      F-Port 21:00:00:e0:8b:04:d5:39
port 1: sw Online      F-Port 21:00:00:e0:8b:04:ee:3a
port 2: sw Online      F-Port 21:00:00:e0:8b:03:8f:e6
port 3: sw Online      F-Port 21:00:00:e0:8b:04:13:3a
port 4: sw Online      F-Port 21:00:00:e0:8b:05:c4:07
port 5: sw Online      F-Port 21:00:00:e0:8b:05:0f:08
port 6: sw Online      F-Port 21:00:00:e0:8b:05:d2:07
port 7: sw Online      F-Port 21:00:00:e0:8b:05:08:08
port 8: sw Online      F-Port 21:00:00:e0:8b:04:e9:39
port 9: sw Online      L-Port 1 public
port 10: sw Online     F-Port 21:00:00:e0:8b:03:97:e6
port 11: sw Online     L-Port 1 public
port 12: sw Online     F-Port 21:00:00:e0:8b:05:d5:07
port 13: sw Online     F-Port 50:05:07:63:00:c3:91:74
port 14: sw Online     F-Port 21:00:00:e0:8b:05:07:08
port 15: sw Online     F-Port 50:05:07:63:00:c7:91:74
```

3. Confirm that each QLogic host adapter and storage unit host adapter has performed a fabric log in to the switch.

Displaying zoning information for the switch

Complete this task to to display zoning information for the switch on an SGI host.

To display the zoning information for the switch, type **cfgShow**. The following is example output. This zoning example is based on the unique worldwide port names of the host adapters and the storage unit fibre-channel ports. This is known as soft zoning. Each path to a storage unit LUN has a separate zone.

```

osvl2109c:admin> cfgShow
Defined configuration:
cfg:  osvl_lab
      dc_1; w2k_1; MS6000Cluster; MSHPCluster; Tivoli; compaqzone1;
      compaqzone2; MS8500Cluster; AIX_ZONE; OSPL3Zone;
      MSCompaqCluster; SGI_Path1; SGI_Path2; NW;
      Netfyzone1; Netfyzone2; Netfyzone3; Netfyzone4

...
zone:  SGI_Path1
      SGI01; osp13_b1a2
zone:  SGI_Path2
      SGI02; osp13_b4a2
...
alias: SGI01    21:00:00:e0:8b:04:d5:39
alias: SGI02    21:00:00:e0:8b:04:ee:3a
alias: osp13_b1a2
      50:05:07:63:00:c3:91:74

alias: osp13_b4a2
      50:05:07:63:00:c7:91:74

...
Effective configuration:
cfg:  osvl_lab
...
zone:  SGI_Path1
      21:00:00:e0:8b:04:d5:39
      50:05:07:63:00:c3:91:74

zone:  SGI_Path2
      21:00:00:e0:8b:04:ee:3a
      50:05:07:63:00:c7:91:74

...

```

Locating the WWPN for an SGI host

Complete this task to locate the WWPN for an SGI host adapter.

1. Restart the server.
2. Type `scsiha -w [bus_number | device]`.

The following is example output:

```

# scsiha -w 6 7 8 9
6 Portname: 210000e08b05d207
7 Portname: 210000e08b04d539
8 Portname: 210000e08b050808
9 Portname: 210000e08b038fe6
#

```

Confirming storage connectivity on a SGI host

Use the procedures in this section to confirm SGI storage connectivity for both the switched fabric and arbitrated loop topologies.

Confirming storage connectivity for SGI in a switched-fabric topology

Confirming storage connectivity for SGI in a switched-fabric topology

This section contains an example of a switched fabric connection with the fibre-channel adapter cards. The following conditions apply to this example:

- The PCI-FC-1P-OPT adapter cards are on PCI busses 6 and 8.
 - There are four LUNs defined on the storage unit.
 - There are two connections from the switch to two host adapters on the storage unit.
 - The switch is not zoned. Therefore, each LUN is available on each host adapter from each host bus adapter. This means that there are four instances of each LUN for each adapter card.
1. Type `scsiha -r {bus_number | device}` at the root of the shell prompt. `bus_number` is the numerical value for the bus and `device` is numerical value of the device.
 2. Type `scsiha -p {bus_number | device}`.
 3. Use the `ioconfig -d -f /hw` command to assign logical controller numbers in the hardware graph to each of the newly discovered physical devices.
 4. Use the `hinv -c disk` command to confirm that all four instances of each of the four LUNs exists.
 5. Use the `ls -l /hw/rdisk` command to verify that you named the directory the same name as the worldwide node name of the storage unit.
 6. Use the `ls -l /hw/rdisk` command to verify that directories for each LUN exist in the `/hw/rdisk/<wwnn>` directory, where `<wwnn>` is the worldwide node name of the device.
 7. Type `cd /hw/rdisk/<wwnn>`, where `<wwnn>` is the worldwide node name of the device, to change to the `/hw/rdisk/<wwnn>` directory. You must execute the next command from within this directory.
 8. Use the `ls -lR lun*` command to verify that the command created the character-special files in each LUN* directory. The asterisk (*) acts as a universal wildcard for any characters that might appear after the initial "LUN" characters.

Confirming storage connectivity for SGI in a fibre-channel arbitrated loop topology

Confirming storage connectivity for SGI in a fibre-channel arbitrated loop topology

This section contains an example of an arbitrated loop connection with the fibre-channel adapter cards. The following conditions apply to this example:

- The PCI-FC-1P-OPT adapter cards are on PCI busses 6 and 8.
 - There are four LUNs defined on the storage unit.
 - Each host adapter is connected to a host adapter in the storage unit.
 - Each host adapter maintains two instances of each LUN.
1. Type `scsiha -r {bus_number | device}` at the root of the shell prompt. `bus_number` is the numerical value for the bus and `device` is the numerical value for the device.
 2. Type `scsiha -p {bus_number | device}`.
 3. Use the `ioconfig -d -f /hw` command to assign logical controller numbers in the hardware graph to each of the newly discovered physical devices.
 4. Use the `hinv -c disk` command to confirm that two instances of each of the four LUNs exist.
 5. Use the `ls -l /hw/disk` command to verify that the command created the appropriate block special files in `hw/disk` for each LUN.

Configuring the storage unit for host failover using IRIX multipathing

This section describes how to configure the storage unit for host failover using IRIX multipathing.

IRIX multipathing does not perform dynamic load-balancing; it only provides path failover. You can achieve static load-balancing if you distribute the primary paths evenly across the controllers when you set up the failover groups in the */etc/failover.conf* file.

When a failover occurs to the alternate path, you must use the */sbin/scsifo -s xxxx* command to manually return to the primary path. Where *xxxx* is the pathname.

When you have a host system with multiple host bus adapters or multiple paths to the storage, IRIX automatically chooses the first path it encounters in the hardware tree as the path to the storage. To prevent this from occurring when you perform a static load-balance, you must configure the *failover.conf* file to point the I/O operations to the correct path. This ensures that you assign a failover path or a backup path for all I/O operations.

Host failover using IRIX multipathing requires the use of the XLV volume manager or in the case of IRIX clusters, the XVM volume manager.

Note: You can also use XVM in non-cluster mode.

Confirming the availability of failover

Complete this task to confirm the availability of failover on an SGI host.

1. Type **chkconfig** to determine if failover is enabled.
2. If failover is not enabled, turn it on using the following example:

```
#chkconfig | grep failover
failover                off
#chkconfig failover on
#/etc/init.d/failover init
#Configuring Failover.
```

3. Start the failover script or restart the server.

Using XVM volume manager for IRIX multipathing failover

To manage host failover for IRIX clusters, you must use the XVM volume manager to configure the *failover.conf* file so that IRIX points I/O operations to the correct path.

Figure 23 on page 159 provides an example of the process that you can use to configure the *failover.conf* file.

```

SGI# /usr/bin/fx -x -d /dev/rdisk/WWNN/lunXvol/PORTNAME
r <press enter> (repartition)
o <press enter> (option disk)
(xfs) <press enter> (xsf filesystem)
y <press enter> (yes)
... <press enter> (previous selection)
exit <press enter>

create /etc/failover.conf

SGI# xvm
xvm: local> label -name NAME/dev/rdisk/WWNN/lunXvol/PORTNAME
xvm: local> slice -volname VOLNAMEphys/NAME
xvm: local> exit

SGI# mkfs -t xfs /dev/lxvm/VOLNAME

SGI# mkdir /DIRECTORY_NAME

SGI# vi /etc/fstab
(for example: /dev/lxvm/VOLNAME/DIRECTORY_NAME xfs rw 0 0)

SGI# mount -a

```

Figure 23. Configuration example for the XVM volume manager

Where

WWNN is the worldwide port name

X is the LUN number

PORTNAME is the port ID

NAME is your assigned nickname for the disk

VOLNAME is the volume ID

DIRECTORY_NAME is the name of the directory where you are storing the failover.conf file

Making a connection through a switched-fabric topology

Complete this task to make a connection through a switched-fabric topology

1. Edit the /etc/failover.conf file to add primary and secondary paths to the storage unit LUNs. “Confirming storage connectivity for SGI in a switched-fabric topology” on page 156 describes the paths of the character-special files. These files begin with the worldwide node name of the storage unit.

The following example shows an edited /etc/failover.conf file.

```

c6lun0      5005076300c09174/lun0/c6p30000 \
            5005076300c09174/lun0/c8p70000

c6lun1      5005076300c09174/lun1/c8p70000 \
            5005076300c09174/lun1/c6p30000

c6lun2      5005076300c09174/lun2/c6p30000 \
            5005076300c09174/lun2/c8p70000

c6lun3      5005076300c09174/lun3/c8p70000 \
            5005076300c09174/lun3/c6p30000

```

2. To confirm the failover configuration, use the IRIX /sbin/scsiifo -d command.

Making a connection through an arbitrated-loop topology

Complete this task to make a connection through an arbitrated-loop topology

1. Edit the `/etc/failover.conf` file to add primary and secondary paths to storage unit LUNs.

The following example shows an edited `/etc/failover.conf` file that adds paths to the storage unit LUNs.

```
GroupA sc3d0111 sc4d0111
GroupB sc3d0117 sc4d0117
```

2. To confirm the proper failover configuration, use the IRIX `/sbin/scsifo -d`.

Switching I/O operations between the primary and secondary paths

Switching I/O operations between the primary and secondary paths

Use the `sbin/scsifo -s` command to manually switch from the primary path to the alternate path. With this command, you can also specify the path from where you want to switch within a failover group. For example, you might want to perform this action to switch back to the initial primary path after a path failure.

Configuring storage in a switched fabric or arbitrated loop topology

The following sections describe how to configure storage in either a switched fabric or arbitrated loop topology.

Note: The `/hw` directory is used to build the hardware graph. The hardware graph represents the collection of all significant hardware that is connected to a system. Do not specify the `/hw` entries in commands that pertain to disk devices. Use the traditional `/dev/rdisk` entries instead. The following example shows a naming convention for the disk devices.

```
#/usr/bin/fx -x -d /dev/rdisk/5005076300c003b4/1un1vo1/c6p30000
```

Configuring storage in a switched fabric topology

This section describes the commands you can use to configure storage in a switched fabric topology

Figure 24 on page 161 and Figure 25 on page 162 shows the standard IRIX storage configuration utilities that you can use to partition and format the storage unit LUNs and to create and mount the file systems. The commands are in bold type.

Note: You must be a root user to use these commands.

```

#/usr/bin/fx -x -d /dev/rdisk/5005076300c003b4/1un1vol/c6p30000
fx version 6.5, Jul 7, 2000
...opening /dev/rdisk/5005076300c003b4/1un1vol/c3p1000

...drive selftest...OK
Scsi drive type == IBM      2105F20      1206

----- please choose one (? for help, .. to quit this menu)-----
[exi]t          [d]ebug/          [l]abel/          [a]uto
[b]adblock/     [ex]ercise/       [r]epartition/

fx> repartition

----- partitions-----
part type      blocks          Megabytes  (base+size)
 7: xlv         4096 + 71798784      2 + 35058
 8: volhdr      0 + 4096             0 + 2
10: volume      0 + 71802880        0 + 35060

capacity is 71802880 blocks

----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive     [o]ptiondrive       [e]xpert
[u]srrootdrive  [re]size

fx/repartition> o

fx/repartition/optiondrive: type of data partition = (xfs) xfs
Warning: you will need to re-install all software and restore user data
from backups after changing the partition layout. Changing partitions
will cause all data on the drive to be lost. Be sure you have the drive
backed up if it contains any user data. Continue? yes

----- partitions-----
part type      blocks          Megabytes  (base+size)
 7: xfs         4096 + 71798784      2 + 35058
 8: volhdr      0 + 4096             0 + 2
10: volume      0 + 71802880        0 + 35060

capacity is 71802880 blocks

----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive     [o]ptiondrive       [e]xpert
[u]srrootdrive  [re]size

fx/repartition> ..

----- please choose one (? for help, .. to quit this menu)-----
[exi]t          [d]ebug/          [l]abel/          [a]uto
[b]adblock/     [ex]ercise/       [r]epartition/

fx>exit

```

Figure 24. Example commands for the IRIX switched fabric storage configuration utility

```

# xlv_make
xlv_make> vol Sharks0
Sharks0
xlv_make> data
Sharks0.data
xlv_make> plex
Sharks0.data.0
xlv_make> ve -force "/dev/dsk/5005076300c003b4/1un1s7/c6p3s0000"
xlv_make> quit

# ls -l /dev/xlv
total 0
brw----- 1 root    sys      192, 10 Mar 12 16:06 Sharks0

# mkfs -t xfs /dev/xlv/Sharks0
meta-data=/dev/xlv/Sharks0    isize=256    agcount=8, agsize=27403 blks
data      =                   bsize=4096    blocks=219218, imaxpct=25
          =                   sunit=0       swidth=0 blks, unwritten=1
naming    =version 1          bsize=4096
log       =internal log      bsize=4096    blocks=1168
realtime  =none              extsz=65536    blocks=0, rtextents=0

# mkdir /lv1_mount
# mount -t xfs /dev/xlv/Sharks0 /lv1_mount
# df -k

```

Filesystem	Type	kbytes	use	avail	%use	Mounted on
/dev/root	xfs	1961580	1750112	211468	90	/
/dev/xlv/Sharks0	xfs	872200	144	872056	1	/lv1_mount

```

#

```

Figure 25. Example commands for the IRIX switched fabric storage configuration utility

Configuring storage in an arbitrated loop topology

This section describes the commands you can use to configure storage in an arbitrated loop topology.

Figure 26 on page 163 and Figure 27 on page 164 shows the standard IRIX storage configuration utilities that you can use to partition and format the storage unit LUNs and to create and mount file systems. The commands are in bold type.


```
# fx -x -d /dev/rdisk/dks3d0117s0

...drive selftest...OK
Scsi drive type == IBM      2105F20      1206

----- please choose one (? for help, .. to quit this menu)-----
[exi]t          [d]ebug/          [l]abel/          [a]uto
[b]adbblock/    [ex]ercise/      [r]epartition/

fx> repartition

----- partitions-----
part type      blocks          Megabytes  (base+size)
 7: xlv        4096 + 71798784      2 + 35058
 8: volhdr     0 + 4096          0 + 2
10: volume     0 + 71802880        0 + 35060

capacity is 71802880 blocks

----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive     [o]ptiondrive    [e]xpert
[u]srrootdrive  [re]size
fx/repartition> o

fx/repartition/optiondrive: type of data partition = (xfs) xfs
Warning: you will need to re-install all software and restore user data
from backups after changing the partition layout. Changing partitions
will cause all data on the drive to be lost. Be sure you have the drive
backed up if it contains any user data. Continue? yes

----- partitions-----
part type      blocks          Megabytes  (base+size)
 7: xfs        4096 + 71798784      2 + 35058
 8: volhdr     0 + 4096          0 + 2
10: volume     0 + 71802880        0 + 35060

capacity is 71802880 blocks
----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive     [o]ptiondrive    [e]xpert
[u]srrootdrive  [re]size
fx/repartition> ..

----- please choose one (? for help, .. to quit this menu)-----
[exi]t          [d]ebug/          [l]abel/          [a]uto
[b]adbblock/    [ex]ercise/      [r]epartition/

fx>exit
```

Figure 26. Example commands for the IRIX arbitrated loop storage configuration utility

```

# xlv_make
xlv_make> vol Sharks0
Sharks0
xlv_make> data
Sharks0.data
xlv_make> plex
Sharks0.data.0
xlv_make> ve -force "/dev/rdisk/dks3d0117s0"
xlv_make> quit

# ls -l /dev/xlv
total 0
brw----- 1 root    sys      192, 10 Mar 12 16:06 Sharks0

# mkfs -t xfs /dev/xlv/Sharks0
meta-data=/dev/xlv/Sharks0    isize=256    agcount=8, agsize=27403 blks
data      =                   bsize=4096    blocks=219218, imaxpct=25
          =                   sunit=0       swidth=0 blks, unwritten=1
naming    =version 1          bsize=4096
log        =internal log      bsize=4096    blocks=1168
realtime   =none              extsz=65536    blocks=0, rtextents=0

# mkdir /lv1_mount
# mount -t xfs /dev/xlv/Sharks0 /lv1_mount
# df -k

```

Filesystem	Type	kbytes	use	avail	%use	Mounted on
/dev/root	xfs	1961580	1750112	211468	90	/
/dev/xlv/Sharks0	xfs	872200	144	872056	1	/lv1_mount

```

#

```

Figure 27. Example commands for the IRIX arbitrated loop storage configuration utility

Sun host attachment

This chapter describes how you attach a Sun host system to a DS6000 with fibre-channel adapters.

This chapter contains the following sections:

- Installing fibre-channel adapters and adapter drivers on a Sun host
- Configuring host device drivers on a Sun host
- Setting system parameters on a Sun host
- Locating the WWPN for a Sun host
- Attaching a Sun host using Storage Traffic Manager System
- Attaching a Sun host using Sun Cluster

Before you attach a host, review the following information:

- The “General requirements for attaching a host system” on page 29
- The DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>
- The latest recommended patch cluster releases that are associated with your operating system at <http://sunsolve.sun.com>. Download and apply if necessary. The recommended patch cluster updates the kernel, SSD driver, ISP driver, and SD driver patches to the current level.

Notes:

1. You can use the SDD for Sun host hardware platforms with 32- or 64-bit mode capabilities that run Solaris 6, 7, 8, 9, or 10.
2. The SDD does not support the Sun host system in a clustering environment.

Installing fibre-channel adapters and adapter drivers on a Sun host

This section describes how you install fibre-channel adapters and adapter drivers on a Sun host.

Installing an AMCC adapter and driver on a Sun host

Complete this task to install an AMCC PCI adapter and driver on a Sun host running Solaris 6, 7, 8 or 9.

Note: AMCC adapters do not support Sun host systems running Solaris 10.

The following steps are general steps used to install AMCC adapters and drivers. Before using these steps to install the adapter and driver, read through all documentation that comes with your adapter card and driver download, and follow any specific instructions included in that documentation.

1. Shut down your host system.
2. Insert the AMCC PCI adapter into an available PCI slot.
3. Connect the cable to the adapter and to the DS6000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS6000, see Host configuration using the DS Storage Manager.
4. Restart the system.
5. Use the instructions in “Downloading and installing a host adapter driver” on page 30 to download and install the AMCC driver for your adapter.

Installing an Emulex adapter and driver on a Sun host

Complete this task to install an Emulex adapter and driver on a Sun host running Solaris 6, 7, 8, 9, and 10.

Note: The Emulex LP9402DC adapter card is a 3.3 volt card only and requires a 3.3 volt PCI slot.

Single- and dual-port, fibre-channel interfaces with an Emulex adapter card support the following public and private loop modes:

- Target
- Public initiator
- Private initiator
- Target and public initiator
- Target and private initiator

To install an Emulex adapter and driver, perform the following steps:

1. Record the IEEE number that is printed on the card. You can use the IEEE number to determine the WWPN.
2. Shut down your host system.
3. Insert the Emulex adapter into an available PCI slot.

4. Connect the cable to the adapter and to the DS6000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS6000, see Host configuration using the DS Storage Manager.
5. Restart the system.
6. Use the instructions in “Downloading and installing a host adapter driver” on page 30 to download and install the Emulex driver for your adapter.

Installing a QLogic adapter and driver on the Sun host

Complete this task to install a QLogic adapter and driver on a Sun host running Solaris 6, 7, 8, 9, and 10.

To install a QLogic adapter and driver, perform the following steps:

1. Shut down your host system.
2. Insert the QLogic adapter into a PCI or SBUS slot on your host.
3. Connect the cable to the adapter and to the DS6000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS6000, see Host configuration using the DS Storage Manager.
4. Restart the system.
5. Use the instructions in “Downloading and installing a host adapter driver” on page 30 to download the QLogic driver for your adapter.
6. Install the QLogic host adapter driver package:

- a. Locate the **qlaxxxx.Z** package image file that you downloaded in step 5.

- b. Log in to your Solaris system as root.

- c. Type **cd/tmp** to change the directory to a temporary directory.

- d. Type the following command to copy the compressed package file to the /tmp directory:

```
cp source_directory/qlaxxxx.Z
```

where **source_directory** is the directory where the compressed package file resides after downloading the driver.

- e. Type the following command to uncompress the package file:

```
uncompress qlaxxxx.Z
```

Note: You must use an upper case **Z** when you specify the file name.

- f. Type the following command to install the uncompressed QLogic driver package:

```
pkgadd -d glaxxxx.Z
```

Note: You must use an upper case **Z** when you specify the file name.

Installing a Sun adapter and driver on a Sun host system

Complete this task to install a Sun-branded adapter and driver on a Sun host running Solaris 6, 7, 8, 9, and 10.

The following steps are general steps to install Sun-branded adapters and drivers. Before you perform these steps to install the adapter and driver, read through all documentation that comes with your adapter card and driver download, and follow any specific instructions that are included in that documentation.

1. Shut down your host system.
2. Insert the Sun adapter into an available PCI slot.

3. Connect the cable to the adapter and to the DS6000 or to the switch port that you have configured for the host adapter. If you have not configured host adapter ports on your DS6000, see Host configuration using the DS Storage Manager.
4. Restart the system.
5. Follow the instructions in “Downloading and installing a host adapter driver” on page 30 to download the latest Sun driver for your adapter.
6. Log in to your Solaris system with root privileges.
7. Navigate to the directory where the compressed package file resides. Type **cd/source_directory**, where **source_directory** is the directory where the compressed package file resides after you have downloaded the driver from the Sun Web site.
8. Type the following command to uncompress the package file:
tar xvf SAN_4[1].4.13_install_it.tar
9. Navigate to the following directory:
cd SAN_4.4.13_install_it
10. Type the following command to install the driver:
./install_it
11. Restart the system.
12. When the system has completed the restart, type the following command to determine the connected fibre-channel adapters:
cfgadm -al
The following output indicates that one fibre-channel adapter was discovered:
c3::500507630304c093 ... unconfigured
13. To configure the adapters so that the Solaris operating system can recognize the logical unit numbers (LUNs), type the following command:
cfgadm -c configure c3::500507630304c093
where **c3::500507630304c093** is the output from step 12 and the WWPN of the DS6000 adapter port that you connected to in step 3.
14. Repeat step 13 as necessary to ensure that all LUNs are recognized.
15. Scan all attached devices using the following command:
devfsadm
16. Type the **format** command to display all LUNs that were discovered by the Solaris operating system.

Configuring host device drivers on a Sun host

Complete this task to update the configuration files for your adapter and the Solaris device driver.

Notes:

1. The Emulex and Sun-Emulex adapters look the same physically. However, Sun-Emulex adapters include Sun firmware, the Sun label, and a part number beginning with SG instead of LP.
2. The QLogic and Sun-QLogic adapters look the same physically. However, Sun-QLogic adapters include Sun firmware, the Sun label, and a part number beginning with SG instead of QLA.

To configure the host device drivers, perform the following steps:

1. Type **cd /kernel/drv** to change directories.
2. Edit parameter settings in your adapter configuration file:

AMCC adapters

Parameters are in the `/kernel/drv/jnicxxx.conf` file, where the value of `xxx` depends on the AMCC model number. For parameter descriptions and recommendations, see “Parameter settings for AMCC adapters on a Sun host” on page 170.

Emulex adapters

Parameters are in the `/kernel/drv/lpfc.conf` file. For parameter descriptions and recommendations, see “Parameter settings for Emulex adapters on a Sun host” on page 172.

Sun-Emulex adapters

Parameters are in the `/kernel/drv/emlxs.conf` file. For parameter descriptions and recommendations, see “Parameter settings for Sun-Emulex adapters on a Sun host” on page 174.

QLogic adapters

Parameters are in the `/kernel/drv/qla2300.conf` file. For parameter descriptions and recommendations, see “Parameter settings for QLogic QLA23xx adapters with driver level 4.03 on a Sun host” on page 177 or “Parameter settings for QLogic QLA23xx and QLA246x adapters with a San Surf configuration on a Sun host” on page 178.

Sun-QLogic adapters

Parameters are in the `/kernel/drv/qlc.conf` file. For parameter descriptions and recommendations, see “Parameter settings for Sun-QLogic adapters on a Sun host” on page 180.

3. Back up the Solaris device driver configuration file, `/kernel/drv/sd.conf`.
4. Edit the `sd.conf` file to add support for the target and LUN pairs that are configured on the host system.

Notes:

- a. Do not change or remove entries in the `sd.conf` file for preexisting devices. Changing or removing entries can cause your system to become inoperable.
- b. Do not add duplicate target and LUN pairs.

Figure 28 on page 169 shows the lines that you must add to the file to access LUNs 0 to 49 on target 0 for fibre-channel connections.

```

name="sd" class="scsi"
    target=0 lun=0;
name="sd" class="scsi"
    target=0 lun=1;
name="sd" class="scsi"
    target=0 lun=2;
name="sd" class="scsi"
    target=0 lun=3;
name="sd" class="scsi"
    target=0 lun=4;
name="sd" class="scsi"
    target=0 lun=5;
name="sd" class="scsi"
    target=0 lun=6;
name="sd" class="scsi"
    target=0 lun=7;
name="sd" class="scsi"
.
.
.
name="sd" class="scsi"
    target=0 lun=48;
name="sd" class="scsi"
    target=0 lun=49;

```

Figure 28. Example of Solaris device driver configuration file entries for fibre-channel connections

Figure 29 shows the start lpfc auto-generated configuration.

Note: You will delete anything that you put within this auto-generated section if you issue the **pkgrm** command to remove the lpfc driver package. You might want to add additional lines to probe for additional LUNs or targets. Delete any lines that represent lpfc targets or LUNs that are not used.

```

name="sd" parent="lpfc" target=0 lun=0;
name="sd" parent="lpfc" target=1 lun=0;
name="sd" parent="lpfc" target=2 lun=0;
name="sd" parent="lpfc" target=3 lun=0;
name="sd" parent="lpfc" target=4 lun=0;
name="sd" parent="lpfc" target=5 lun=0;
name="sd" parent="lpfc" target=6 lun=0;
name="sd" parent="lpfc" target=7 lun=0;
name="sd" parent="lpfc" target=8 lun=0;
name="sd" parent="lpfc" target=9 lun=0;
name="sd" parent="lpfc" target=10 lun=0;
name="sd" parent="lpfc" target=11 lun=0;
name="sd" parent="lpfc" target=12 lun=0;
name="sd" parent="lpfc" target=13 lun=0;
name="sd" parent="lpfc" target=14 lun=0;
name="sd" parent="lpfc" target=15 lun=0;
name="sd" parent="lpfc" target=16 lun=0;
name="sd" parent="lpfc" target=17 lun=0;
name="sd" parent="lpfc" target=17 lun=1;
name="sd" parent="lpfc" target=17 lun=2;
name="sd" parent="lpfc" target=17 lun=3;

```

Figure 29. Example of a start lpfc auto-generated configuration for a Sun host

5. From the Sun UNIX prompt, type:
 - a. **cd /**
 - b. **touch reconfigure**
 - c. **reboot**

The fibre-channel adapters that are supported for attaching the storage unit to a Sun host are capable of full-fabric support. Ensure that all fibre-channel driver

configurations include worldwide port name, worldwide node name, port ID, or host adapter binding of target LUN pairs.

Binding of target LUN pairs implements the Solaris fibre-channel host adapter configuration file that is installed by the adapter software package. Refer to the manufacturer's adapter documentation and utilities for detailed configuration instructions.

You can tune fibre-channel host adapter configuration files for host system reliability and performance.

Parameter settings for AMCC adapters on a Sun host

You can use these recommended settings for AMCC FCX-6562, AMCC FCX2-6562, AMCC FCE-6460, or AMCC FCE-1473 adapters on a Sun host.

Table 14 describes the configuration parameters and recommended settings for AMCC adapters.

Table 14. Recommended configuration settings for AMCC adapters on a Sun host

Parameters	Recommended settings
BusyRetryDelay	5000 (Default) Specifies the delay, in milliseconds, before retrying after receipt of an I/O operation with a SCSI Busy status from a target. The number of retries is based on the Solaris retry count associated with the I/O operation.
FailoverDelay	30 (Default) Specifies the delay, in seconds, before failing all I/O operations for an offline target. If the delay timer expires, all I/O operations for the failed target is returned to the application. A zero value disables failover.
FcEmldEngTcbCount	1789 (Default) Specifies the total number of concurrent exchanges (also called transfer control blocks) that the adapter allows. To optimize performance, set this parameter to match the memory capacity of the hardware.
FcEngHeartbeatInterval	5 (Default) Specifies the interval used by the adapter when it sends a test frame. When the AMCC adapter or driver detects that the fibre-channel link is up (and there is no I/O activity), it sends a test frame (or heartbeat) to itself to verify link integrity. The test frame is sent at the interval specified by this parameter. If the test frame does not complete, there might be a link problem. In this situation, the driver initiates error recovery to reestablish a good link. A value of 0 disables the heartbeat.
FcFlogiTimeout	10 (Default) Specifies the amount of time (in seconds) that the driver waits for a fabric login (FLOGI) to be accepted. The value should be increased only if the switch to which the host adapter is connected requires more time to respond to a FLOGI. The number of retries is configured with the FlogiRetryCount parameter.

Table 14. Recommended configuration settings for AMCC adapters on a Sun host (continued)

Parameters	Recommended settings
FcLinkSpeed	<p>3 (Default)</p> <p>Specifies the desired fibre-channel link speed:</p> <ul style="list-style-type: none"> 0: Default to EEPROM setting 1: Force 1 gigabit per second 2: Force 2 gigabit per second 3: Auto negotiate link speed
FcLinkUpRecoveryTime	<p>1000 (Default)</p> <p>Specifies the delay, in milliseconds, after the link is up before port discovery begins, allowing the link to stabilize and protecting against a possible I/O surge. This timer is reset every time the link comes up. The default value is adequate for most configurations.</p>
FlogiRetryCount	<p>3 (Default)</p> <p>Specifies the total number of fabric login (FLOGI) attempts before giving up logging in to a switch. Failure prevents participation on a fabric topology.</p>
JniCreationDelay	<p>5 (Default)</p> <p>Specifies the delay, in seconds, after driver creation to allow the network to stabilize, discover ports, and build the driver database. Increase this value if targets are being discovered too late in the restart process.</p>
LunRecoveryInterval	<p>50 (Default)</p> <p>Sets the LUN I/O recovery interval, in milliseconds, after the driver reconnects to a disk. It is a global parameter affecting all targets, and determines how long the driver waits after a port is discovered until sending I/O operations to that port. Some devices might require more time to flush an I/O operation that was in progress prior to a link going down; if this is the case, increase the value of this parameter.</p>
PlogiControlSeconds	<p>30 (Default)</p> <p>Specifies the number of seconds that the driver waits for a successful port login (PLOGI) attempt. The maximum number of attempts is defined by the PlogiRetryCount parameter. Some devices might take longer to respond to PLOGIs; if this is the case, increase the value of this parameter.</p>
PlogiRetryCount	<p>5 (Default)</p> <p>Specifies the total number of port login (PLOGI) attempts before failing to log in to an SCSI target.</p>
QfullRetryCount	<p>5 (Default)</p> <p>Specifies the number of times an I/O operation is retried due to receipt of a SCSI queue full status from a target. The delay between retries is based on the QfullRetryDelay parameter.</p>

Table 14. Recommended configuration settings for AMCC adapters on a Sun host (continued)

Parameters	Recommended settings
QfullRetryDelay	5000 (Default) Specifies the delay, in milliseconds, before retrying after receipt of an I/O operation with a SCSI queue full status from a target. The number of retries is based on the QfullRetryCount parameter.
TargetOfflineEnable	0 Enables or disables target offline when target port fails: 0: Disabled 1: Enabled (Default)
TimeoutResetEnable	0 (Disabled) Specifies whether SCSI target resets for timed-out I/O operations are enabled or disabled. When the timer expires (usually 60 seconds, as specified by the upper layers), the driver issues a target reset to attempt to clear the device (which might be either too busy to respond or hung).

Parameter settings for Emulex adapters on a Sun host

You can use these recommended settings for Emulex LP9002DC, LP9002L, LP9002S, LP9402DC, LP9802, LP10000, LP10000, LP11000, and LP11002 DC adapters.

Table 15 describes the configuration parameters and recommended settings for the Emulex adapters.

Table 15. Recommended configuration settings for Emulex adapters on a Sun host

Parameters	Recommended settings
automap	1 (Default) SCSI IDs for all FCP nodes without persistent bindings are automatically generated. If new FCP devices are added to the network when the system is down, there is no guarantee that these SCSI IDs will remain the same when the system is restarted. If one of the FCP binding methods is specified, automap devices use the same mapping method to preserve SCSI IDs between link down and link up. If no bindings are specified: 1: Forces WWNN binding. 2: Forces WWPN binding. 3: Forces DID binding. If automap is 0, only devices with persistent bindings are recognized by the system.
fcp-on	1 (Default. FCP is on.)

Table 15. Recommended configuration settings for Emulex adapters on a Sun host (continued)

Parameters	Recommended settings
lun-queue-depth	30 Specifies the default value that the driver uses to limit the number of outstanding commands per FCP LUN. This value is global, affecting each LUN recognized by the driver, but can be overridden on a per-LUN basis. You might have to configure RAID using the per-LUN tunable throttles.
no-device-delay	1 Specifies the device delay: 0: No delay (Default) 1: Short delay 2: Long delay Setting a long delay might permit I/O operations to build up, each with a pending timeout, which can result in the exhaustion of critical Solaris kernel resources. In this case, you might see a fatal message such as PANIC: Timeout table overflow.
nodev-tmo	30 Specifies the default value in seconds that the driver holds all I/O errors on devices that do not respond until the timer expires. This value helps the driver avoid transient device errors.
network-on	0 (Default. Recommended for fabric.) 0: IP networking is turned off. 1: IP networking is turned on.
scan-down	0 0: The driver uses an inverted ALPA map, effectively scanning ALPAs from high to low as specified in the FC-AL annex. 1: The driver uses arbitrated loop topology.
tgt-queue-depth	0 Specifies the default value that the driver uses to limit the number of outstanding commands per FCP target. This value is global, affecting each target that is recognized by the driver, but it can be overridden on a per-target basis. You might have to configure RAID using the per-target tunable throttles.
topology	2 (for fabric); 4 (for non-fabric) 2: Point-to-point topology only 4: Arbitrated-loop topology only
xmt-que-size	256 (Default) Specifies the size of the transmit queue for mbufs. You can specify a value of 128 - 10240.
zone-rscn	0 (Default); 1 , if you use soft zoning with Brocade fabrics. Setting this parameter to 1 causes the driver to check with the name server to see if an N_Port ID that is received from an RSCN applies.

Parameter settings for Sun-Emulex adapters on a Sun host

You can use these recommended settings for Sun-Emulex SG-XPCI1/2FC-EMx adapters on a Sun host.

Table 16 describes configuration parameters and recommended settings for Sun-Emulex adapters.

Table 16. Recommended configuration settings for Sun-Emulex adapters on a Sun host

Parameters	Recommended settings
Buffer parameters	
num-iocbs	1024 (Default) Specifies the number of iocb buffers to allocate. You can specify a value of 128 - 10240.
ub-bufs	1000 (Default) Specifies the number of unsolicited buffers to allocate. You can specify a value of 40 - 16320.
Fibre-channel connection parameters	
ack0	0 (Default) Specifies whether ACK1 or ACK is used for class 2 acknowledgement: 0: The driver uses ACK1. 1: The driver uses ACK0.
adisc-support	1 (Default) Specifies the level of driver support for the fibre-channel ADISC login: 0: No support. Flush active I/O for all FCP target devices at link down. 1: Partial support. Flush I/O for non-FCP2 target devices at link down. 2: Full support. Hold active I/O for all devices at link down.
assign-alpha	0x00 (Default) Specifies the preferred ALPA for the adapter. This parameter is valid only if the topology parameter is set to loop mode (4). You can specify valid ALPA values of 0x00 – 0xef. If multiple adapter instances on the same host are on the same loop, you must specify a different value for each adapter, for example: emlxs0-assign-alpha=0x01; assigns ALPA 0x01 to adapter 0 emlxs1-assign-alpha=0x02; assigns ALPA 0x02 to adapter 1 emlxs2-assign-alpha=0x04; assigns ALPA 0x04 to adapter 2
cr-count	1 (Default) This parameter is valid only if the cr-delay parameter enables the coalesce response feature. Specifies the number of I/O completions after which an interrupt response is generated. You can specify a value of 1 - 255.

Table 16. Recommended configuration settings for Sun-Emulex adapters on a Sun host (continued)

Parameters	Recommended settings
cr-delay	<p>0 (Default. The coalesce response feature is disabled.)</p> <p>If the coalesce response feature is enabled, specifies the number of milliseconds after which an interrupt response is generated if the cr-count value has not been satisfied. To enable this feature, specify a value of 1 - 63.</p>
link-speed	<p>0 (Default)</p> <p>Specifies the link speed to be used for initializing the fibre-channel connection:</p> <ul style="list-style-type: none"> 0: Select automatically 1: 1 Gigabaud 2: 2 Gigabauds 4: 4 Gigabauds
num-nodes	<p>512 (Default)</p> <p>Specifies the number of fibre-channel nodes (N_Ports). You can specify a value of 2 - 512.</p>
pm-support	<p>1 (Default)</p> <p>Enables or disables power management support in the driver:</p> <ul style="list-style-type: none"> 0: Support is disabled. 1: Support is enabled.
topology	<p>0 (Default, unless you want to run as an N_Port or an NL_Port.)</p> <p>Specifies the link topology to be used for initializing the fibre-channel connection:</p> <ul style="list-style-type: none"> 0: Attempt loop mode; if fails, attempt point-to-point mode. 2: Attempt point-to-point mode only. Use this setting if you want to run as an N_Port. 4: Attempt loop mode only. Use this setting if you want to run as an NL_Port. 6: Attempt point-to-point mode; if fails, attempt loop mode.
IP networking support parameter	
network-on	<p>1 (Default)</p> <p>Enables or disables IP networking support in the driver:</p> <ul style="list-style-type: none"> 0: Support is disabled. 1: Support is enabled.
Message parameters	
console-errors	<p>0 (Default)</p> <p>Specifies the verbose level for driver errors to the console:</p> <ul style="list-style-type: none"> 0: Sets the minimum level. 0xffffffff: Sets the maximum level.

Table 16. Recommended configuration settings for Sun-Emulex adapters on a Sun host (continued)

Parameters	Recommended settings
console-notice	<p>0 (Default)</p> <p>Specifies the verbose level for driver notices to the console:</p> <p>0: Sets the minimum level.</p> <p>0xffffffff: Sets the maximum level.</p>
console-warnings	<p>0 (Default)</p> <p>Specifies the verbose level for driver warnings to the console:</p> <p>0: Sets the minimum level.</p> <p>0xffffffff: Sets the maximum level.</p>
log-errors	<p>0xffffffff (Default)</p> <p>This parameter sets the verbose level for driver errors to the system log file:</p> <p>0: Sets the minimum level.</p> <p>0xffffffff: Sets the maximum level.</p> <p>To configure the types of events to be logged, see "System log file parameters."</p>
log-notice	<p>0xffffffff (Default)</p> <p>Specifies the verbose level for driver notices to the system log file:</p> <p>0: Sets the minimum level.</p> <p>0xffffffff: Sets the maximum level.</p> <p>To configure the types of events to be logged, see "System log file parameters."</p>
log-warnings	<p>0xffffffff (Default)</p> <p>Specifies the verbose level for driver warnings to the system log file:</p> <p>0: Sets the minimum level.</p> <p>0xffffffff: Sets the maximum level.</p> <p>To configure the types of events to be logged, see "System log file parameters."</p>

System log file parameters

The system log file is usually located in the /var/adm/messages directory. Each parameter is a bit mask that enables or disables messages for specific types of events. If the bit is set, the messages of that type are enabled.

Table 17 lists the system log parameters and settings you can use to specify the types of messages to be logged.

Table 17. System log parameters for configuring Sun-Emulex adapter messages

Event Types	Parameters	Settings
Driver attach and detach events	LOG_DRIVER	0x00000002
ELS events	LOG_ELS	0x00000100

Table 17. System log parameters for configuring Sun-Emulex adapter messages (continued)

Event Types	Parameters	Settings
FCP target mode events	LOG_TGTM	0x00000800
FCP traffic events	LOG_FCP	0x00000400
Firmware download events	LOG_FIRMWARE	0x00008000
General I/O packet events	LOG_PKT	0x00000200
HBA initialization events	LOG_INIT	0x00000004
IP traffic events	LOG_IP	0x00001000
IOCTL events	LOG_IOCTL	0x00004000
Link events	LOG_LINK	0x00000080
Mailbox events	LOG_MBOX	0x00000020
Memory management event	LOG_MEM	0x00000008
Miscellaneous events	LOG_MISC	0x00000001
Node events	LOG_NODE	0x00000040
Service Level Interface events	LOG_SLI	0x00000010
SFS events	LOG_SFS	0x00002000

Parameter settings for QLogic QLA23xx adapters with driver level 4.03 on a Sun host

You can use these recommended configuration settings for QLogic QLA2310F, QLA2340, and QLA2342 adapters with driver level 4.03 on a Sun host.

Table 18 describes the configuration parameters and recommended settings for QLogic QLA23xx adapters with driver level 4.03. These settings are for a host that is attached directly to DS6000 or through a fabric switch.

Table 18. Recommended configuration settings for QLogic QLA23xx adapters with driver level 4.03 on a Sun host

Parameters	Recommended settings
hba0-max-frame-length	=2048;
hba0-max-iocb-allocation	=256;
hba0-execution-throttle	=31;
hba0-login-timeout	=4;
hba0-login-retry-count	=1;
hba0-fabric-retry-count	=10;
hba0-adapter-hard-loop-ID	=0;
hba0-enable or disable-64bit PCI DMA addressing	=0: disabled =1: enabled
hba0-enable or /disable-LIP-reset	=0: disabled
hba0-enable or disable-LIP-full-login	=1: enabled
hba0-enable or disable-target-reset	=0: disabled
hba0-reset-delay	=5;
hba0-port-down-retry-count	=30;
hba0-link-down-error	=1: enabled

Table 18. Recommended configuration settings for QLogic QLA23xx adapters with driver level 4.03 on a Sun host (continued)

Parameters	Recommended settings
hba0-link-down-timeout	=60;
hba0-connection-options	=0: loop only (for direct attachment) =1: point-to-point only (for fabric attachment)
hba0-device-configuration-mode	=1: use port name
hba0-fc-tape	=0: disabled =1: enabled
hba0-fc-data-rate	=2: Auto-negotiate
hba0-command-completion-option	=0: response queue
persistent binding only option	=0: reports the discovery of persistent binded and non-bindeds devices to the operating system

Parameter settings for QLogic QLA23xx and QLA246x adapters with a San Surf configuration on a Sun host

You can use these recommended settings for QLogic QLA23xx and QLogic QLA226x adapters on a Sun host. These adapters are for storage units with a San Surf configuration (4.06+ driver).

Table 19 describes the configuration parameters and recommended settings. These settings are for a host that is attached directly to DS6000 or through a fabric switch.

Table 19. Recommended configuration settings for QLogic QLA23xx and QLogic QLA246x adapters with a San Surf configuration on a Sun host

Parameters	Recommended settings
Host Parameters	
Execution Throttle	16 (Default)
Frame Length	2048
Port Down Retry Count	30
Connection Options	<ul style="list-style-type: none"> • Loop Only • Loop Retrieved • Point-to-Point
Enable Fibre Channel Tape Support	Enable
Enable Fibre Channel Confirm	Enable
Data Rate	2 Auto Negotiate
Advanced Host Parameters (See Note 1)	
Persistent Plus New	Enable
Persistent Target Only	Disable
Persistent Bind HBA	Disable
Fast Error Reporting	Disable
Link Down Error	Enable

Table 19. Recommended configuration settings for QLogic QLA23xx and QLogic QLA246x adapters with a San Surf configuration on a Sun host (continued)

Parameters	Recommended settings
Extended Logging	Disable
Enable Target Reset	Enable
Enable Fast Command Posting	Disable
Enable 4Gb Addressing	Disable
Maximum LUNs per Target	8 (Default)
Loop Down Timeout	60 (Default)
Advanced Firmware Parameters (See Note 1)	
Enable Adapter Hard Loop ID	Disable
Enable LIP Full Login	Enable
Enable Class 2 Service	Disable
Enable ACKO	Disable
Enable Read Transfer Ready	Disable
Hard Loop ID	Default
Login Retry Count	8
Interrupt Delay Timer	0
Wizard (for binding) (See Note 2)	
Notes: <ol style="list-style-type: none"> 1. All of the parameters in the Advanced Host Parameters and Advanced Firmware Parameters categories use the default parameter settings. 2. San Surf uses the Wizard parameter to update the binding configuration information in the qlaxxxx.conf file. Figure 30 shows an example of a qlaxxxx.conf file that contains binding inserts at the end of the file. 	

```

hba3-SCSI-target-id-0-fibre-channel-node-name="5005076300c0863a";
hba3-SCSI-target-id-0-fibre-channel-port-name="5005076300cf863a";

hba3-SCSI-target-id-1-fibre-channel-node-name="5005076300c0863a";
hba3-SCSI-target-id-1-fibre-channel-port-name="5005076300cb863a";

hba2-SCSI-target-id-0-fibre-channel-node-name="5005076300c0863a";
hba2-SCSI-target-id-0-fibre-channel-port-name="5005076300cf863a";

hba2-SCSI-target-id-1-fibre-channel-node-name="5005076300c0863a";
hba2-SCSI-target-id-1-fibre-channel-port-name="5005076300cb863a";

hba1-SCSI-target-id-0-fibre-channel-node-name="5005076300c0863a";
hba1-SCSI-target-id-0-fibre-channel-port-name="5005076300c1863a";

hba1-SCSI-target-id-1-fibre-channel-node-name="5005076300c0863a";
hba1-SCSI-target-id-1-fibre-channel-port-name="5005076300c5863a";

hba0-SCSI-target-id-0-fibre-channel-node-name="5005076300c0863a";
hba0-SCSI-target-id-0-fibre-channel-port-name="5005076300c5863a";

hba0-SCSI-target-id-1-fibre-channel-node-name="5005076300c0863a";
hba0-SCSI-target-id-1-fibre-channel-port-name="5005076300c1863a";

```

Figure 30. Example of a qlaxxxx.conf file for a Sun host

Parameter settings for Sun-QLogic adapters on a Sun host

You can use these recommended configuration settings for Sun-QLogic SG-XPCI1/2FC-QFx adapters on a Sun host.

Table 20 describes configuration parameters and recommended settings for Sun-QLogic adapters. These settings are for a host that is attached directly to DS6000 or through a fabric switch.

Table 20. Recommended configuration settings for Sun-QLogic adapters on a Sun host

Parameters	Recommended settings
adapter-hard-loop-ID	This parameter is valid only when adapter hard loop ID is enabled (enable-adapter-hard-loop-ID=0;). You can specify a value of 0 - 125.
connection-options	<p>When you are connecting the HBA to a switch port with auto-mode sensing capability (a G port), it is recommended that this parameter be set to either 0 or 1. Setting this parameter to 2 when connected to a G port on a switch, can result in device-lost or system-hang conditions. For other configurations, the default value is recommended.</p> <p>Specifies the connection mode that the driver firmware uses:</p> <ul style="list-style-type: none">0: Loop only1: Point-to-point only2: Loop preferred, otherwise point-to-point (default)
enable-adapter-hard-loop-ID	<p>0 (Default)</p> <p>Enables or disables the setting of the hard loop ID for adapters on the fibre-channel bus:</p> <ul style="list-style-type: none">0: Enabled1: Disabled <p>The default value is recommended and might be necessary on some bus configurations where devices fail to appear.</p>
enable-FCP-2-error-recovery	<p>1 (Default)</p> <p>Enables or disables the fibre-channel FCP 2 error recovery support:</p> <ul style="list-style-type: none">0: Disabled1: Enabled
enable-link-down-error	<p>1 (Default)</p> <p>Enables or disables driver error reporting during a link-down condition:</p> <ul style="list-style-type: none">0: Disabled1: Enabled
enable-LIP-full-login-on-bus-reset	<p>1 (Default)</p> <p>Enables or disables the adapter to issue a LIP full login reset during fibre-channel reset:</p> <ul style="list-style-type: none">0: Disabled1: Enabled

Table 20. Recommended configuration settings for Sun-QLogic adapters on a Sun host (continued)

Parameters	Recommended settings
enable-LIP-reset-on-bus-reset	<p>0 (Default)</p> <p>Enables or disables the adapter to issue a LIP reset during fibre-channel reset:</p> <p>0: Disabled 1: Enabled</p>
enable-target-reset-on-bus-reset	<p>0 (Default)</p> <p>Enables or disables the adapter to issue a target reset during fibre-channel reset:</p> <p>0: Disabled 1: Enabled</p> <p>The default value is recommended.</p>
execution-throttle	<p>32 (Default)</p> <p>Specifies the maximum number of commands sent per LUN by the ISP2xxx firmware. You can specify a value of 1 - 65535.</p> <p>Note: Exceeding device capabilities causes unnecessary command retries that impact performance.</p>
extended-logging	<p>0 (Default)</p> <p>Enables or disables the logging of driver-detected events that occur in the driver or fibre-channel bus:</p> <p>0: Disabled 1: Enabled</p> <p>Events are usually logged in the /var/adm/messages file.</p>
fc-data-rate	<p>2 (Default)</p> <p>Specifies the fibre-channel data rate which the driver uses:</p> <p>0: 1 gigabit/second 1: 2 gigabits/second 2: Automatically negotiate 3: 4 gigabits/second</p>
firmware-dump-flags	<p>0 (Default)</p> <p>Specifies whether to force a firmware dump when one or more flag conditions occur:</p> <p>0: Disabled 1: Enabled</p> <p>Flag conditions:</p> <p>1: TAKE_FW_DUMP_ON_MAILBOX_TIMEOUT 2: TAKE_FW_DUMP_ON_ISP_SYSTEM_ERROR 4: TAKE_FW_DUMP_ON_DRIVER_COMMAND_TIMEOUT 8: TAKE_FW_DUMP_ON_LOOP_OFFLINE_TIMEOUT</p>
link-down-timeout	<p>0 (Default)</p> <p>Specifies the number of seconds that the driver waits for a fibre-channel loop to come up before reporting the failure. You can specify a value of 0 - 240.</p>

Table 20. Recommended configuration settings for Sun-QLogic adapters on a Sun host (continued)

Parameters	Recommended settings
login-retry	<p>4 (Default)</p> <p>Specifies the maximum number of retry attempts for the firmware when the login into a device fails. You can specify a value of 0 - 255.</p>
login-timeout	<p>3 (Default)</p> <p>Specifies the maximum number of seconds that the driver waits for the login of a device to complete. You can specify a value of 0 - 255.</p> <p>Note: Large values can cause long delays during initialization and device reconfiguration.</p>
max-frame-length	<p>2048 (Default)</p> <p>Specifies the frame payload length, in bytes, that is used by the ISP2xxx firmware. Other possible values are 1024 and 2048.</p>
pci-max-read-request	<p>2048 (Default)</p> <p>Specifies the number of bytes to be used for the PCI maximum read request, and overrides the value programmed by the system. Other possible values are 128, 256, 512, 1024, and 4096.</p> <p>Note: If this value is not specified or is not equal to 128, 256, 512, 1024, 2048 or 4096, the ISP2xxx defaults to the value specified by the system.</p>
port-down-retry-count	<p>8 (Default)</p> <p>Specifies the number of command retries when devices are not responding on the fibre-channel bus. You can specify a value of 0 - 255.</p> <p>Note: Large values can cause long delays for failover software to detect a failing device.</p>
port-down-retry-delay	<p>2 (Default)</p> <p>Specifies the number of seconds to delay between port-down retries. You can specify a value of 0 - 255.</p> <p>Note: Large values can cause long delays for failover software to detect a failing device, and can also cause upper layer drivers or applications to time out IO operations.</p>
queue-full-retry-count	<p>16 (Default)</p> <p>Specifies the number of times to retry an SCSI queue-full error. You can specify a value of 0 - 255.</p>
queue-full-retry-delay	<p>2 (Default)</p> <p>Specifies the number of seconds to delay, after an SCSI queue-full error, before starting any new I/O commands. You can specify a value of 0 - 255.</p>
reset-delay	<p>5 (Default)</p> <p>Specifies the number of seconds to delay after a reset before sending commands to the devices on the fibre-channel bus. You can specify a value of 1 - 255.</p>

Setting system parameters on a Sun host

This section describes how you set system parameters for AMCC, Emulex, QLogic, Sun-Emulex, and Sun-QLogic adapters on a Sun host.

If you do not set a parameter in the system file, the default value will be used if the parameter exists in the operating system level. For example, the default value for the `ssd_retry_count` parameter will not be used for Solaris 10 because this parameter is not valid for Solaris 10.

This section does not describe all system parameters and their default values. To find out the default value for any system parameter, type the following command:

```
echo 'system_parameter_name/D' |mdb -k
```

For example:

```
echo 'ssd_min_throttle/D' |mdb -k
```

The following is the command output:

```
ssd_min_throttle:  
ssd_min_throttle: 8
```

Setting system parameters for AMCC adapters on a Sun host

Complete this task to set the system parameters for AMCC adapters on a Sun host.

To set the system parameters, perform the following steps:

1. Type `cd/etc` to change to the `/etc` subdirectory.
2. Back up the system file in the subdirectory.
3. Edit the system file and set the following parameters for servers with configurations that use AMCC adapters:

Note: To configure the AMCC adapters see the EZ Fibre Configuration Utility at <http://www.amcc.com>.

sd_max_throttle

This parameter specifies the maximum number of commands that the `sd` driver can queue to the host adapter driver. The default value is 256, but you must set the parameter to a value less than or equal to a maximum queue depth for each LUN connected. Determine the value by using the following formula: $256 \div (\text{LUNs per adapter})$ where LUNs per adapter is the largest number of LUNs assigned to a single adapter.

To set the `sd_max_throttle` parameter for the storage unit LUNs in this example, add the following line to the `/etc/system` file:

```
set sd:sd_max_throttle=5
```

sd_io_time

This parameter specifies the timeout value for disk operations. Add the following line to the `/etc/system` file to set the `sd_io_time` parameter for the storage unit LUNs:

```
set sd:sd_io_time=0x78
```

sd_retry_count

This parameter specifies the retry count for disk operations. Add the following line to the `/etc/system` file to set the `sd_retry_count` parameter for the storage unit LUNs:

```
set sd:sd_retry_count=5
```

maxphys

This parameter specifies the maximum number of bytes that you can transfer for each SCSI transaction. The default value is 126976 (124 KB). If the I/O block size that you requested exceeds the default value, the request is broken into more than one request. You must tune the value for the application requirements. For maximum bandwidth, set the maxphys parameter by adding the following line to the /etc/system file:

```
set maxphys=1048576
```

Note: Do not set the value for maxphys greater than 1048576 (1 MB). Setting this value too high can cause the system to hang.

Setting system parameters for Emulex or QLogic adapters on a Sun host

Complete this task to set the system parameters for Emulex or QLogic adapters on a Sun host.

To set the system parameters, perform the following steps:

1. Type **cd /etc** to change to the /etc subdirectory.
2. Back up the system file in the subdirectory.
3. Edit the system file and set the following parameters for servers with configurations that only use Emulex or QLogic adapters:

sd_max_throttle

This parameter specifies the maximum number of commands that the sd driver can queue to the host adapter driver. The default value is 256, but you must set the parameter to a value less than or equal to a maximum queue depth for each connected LUN. Determine the value by using the following formula: $256 \div (\text{LUNs per adapter})$ where LUNs per adapter is the largest number of LUNs that are assigned to a single adapter.

To set the sd_max_throttle parameter for the storage unit LUNs in this example, add the following line to the /etc/system file:

```
set sd:sd_max_throttle=5
```

sd_io_time

This parameter specifies the timeout value for disk operations. Add the following line to the /etc/system file to set the sd_io_time parameter for the storage unit LUNs:

```
set sd:sd_io_time=0x78
```

sd_retry_count

This parameter specifies the retry count for disk operations. Add the following line to the /etc/system file to set the sd_retry_count parameter for the storage unit LUNs:

```
set sd:sd_retry_count=5
```

maxphys

This parameter specifies the maximum number of bytes you can transfer for each SCSI transaction. The default value is 126976 (124 KB). If the I/O block size that you requested exceeds the default value, the request is broken into more than one request. You must tune the value to the application requirements. For maximum bandwidth, set the maxphys parameter by adding the following line to the /etc/system file:

```
set maxphys=1048576
```

Note: Do not set the value for maxphys greater than 1048576 (1 MB). Setting this value too high can cause the system to hang.

Setting system parameters for Sun-Emulex or Sun-QLogic adapters on a Sun host

Complete this task to set the system parameters for Sun-Emulex or Sun-QLogic adapters on a Sun host.

To set the system parameters, perform the following steps:

1. Type `cd /etc` to change to the `/etc` subdirectory.
2. Back up the system file in the subdirectory.
3. Edit the system file and set the following parameters for servers with configurations that only use Sun-branded host adapters:

ssd_io_time

This parameter specifies the timeout value for disk operations. Add the following line to the `/etc/system` file to set the `ssd_io_time` parameter for the storage unit LUNs:

```
set ssd:ssd_io_time=0x78
```

ssd_max_throttle

This parameter specifies the maximum number of commands that the SSD driver can queue to the host adapter driver. The default value is 256, but you must set the parameter to a value less than or equal to a maximum queue depth for each connected LUN. Determine the value by using the following formula: $256 \div (\text{LUNs per adapter})$ where LUNs per adapter is the largest number of LUNs that are assigned to a single adapter.

To set the `ssd_max_throttle` parameter for the storage unit LUNs in this example, add the following line to the `/etc/system` file:

```
set ssd:ssd_max_throttle=5
```

Notes:

- a. Theoretically, all LUNs on a port must be simultaneously busy to reach the maximum throttle setting that is specified on the `ssd_max_throttle` parameter times the number of LUNs. You can tune the `ssd_max_throttle` setting higher than the formula dictates; however, this might result in an unexpected workload spike that can cause queue overrun.
- b. In general, an `ssd_max_throttle` setting of less than 8 can have a negative performance impact for some types of workloads that benefit from queueing. Generally, you can avoid the requirement for low throttle settings by limiting the number of LUNs that you configure on a port.

Locating the WWPN for a Sun host

Complete this task to locate the WWPN for an AMCC, Emulex, QLogic, Sun-Emulex, or Sun-QLogic adapter on a Sun host.

Note: If you have installed multiple host adapters, you will see more than one WWPN.

To locate the WWPN, perform the following steps:

1. After you install the adapter and restart the host system, open the `/var/adm/messages` file.
2. Search for the line that contains the WWPNN.
 - AMCC adapters:
 - For the 2Gb AMCC adapter, search for `AMCCc146xN: Fibre Channel WWPNN`, where *N* is the adapter number (0, 1, and so on). You can find the WWPNN on the same line following the WWPNN.
 - For the AMCC PCI adapter, search for `fca-pcix: Fibre Channel WWPNN`, where *x* is the adapter number (0, 1, and so on). You can find the WWPNN on the same line following the WWPNN.
 - For the AMCC SBUS adapter, search for `fcawx: Fibre Channel WWPNN`, where *x* is the adapter number (0, 1, and so on). You can find the WWPNN on the same line immediately after the WWPNN.
 - For the Emulex adapter, search for `lpfcx: Fibre Channel WWPNN`, where *x* is the adapter number 0, 1, and so on).
 - For the QLogic QLA23xx adapter, search for `qla23200-hbax-adapter-port-name` where *x* is the adapter number (0, 1, and so on).
 - For the Sun-Emulex adapter, search for `emlxx: Fibre Channel WWPNN`, where *x* is the adapter number 0, 1, and so on).
 - For the Sun-QLogic adapter, search for `qlcx: Fibre Channel WWPNN`, where *x* is the adapter number 0, 1, and so on).

Attaching a Sun host using Storage Traffic Manager System

This section describes how you attach a Sun host to a DS6000 using Storage Traffic Manager System.

The Storage Traffic Manager System (STMS) is a multipathing mechanism for native Sun systems. Using the DS CLI, you can specify **Sun - Solaris** as the port profile. Using the Storage Manager GUI, you can specify **Sun servers (Solaris) (Sun)** as the host system when attaching the storage unit to a host operating with STMS.

See the Sun Web site at <http://www.sun.com> for the most current documentation.

Notes:

Use the following Sun documentation for the required configuration changes to enable Sun MPxIO multipathing software with a DS6000.

- For Solaris 8 and 9, see *Sun StorEdge Traffic Manager Installation and Configuration Guide For the Solaris OS & Sun SAN Foundation Software 4.4* at <http://docs.sun.com/app/docs/doc/817-3674-12?q=817-3674>.
- For Solaris 10, see *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide* at <http://docs.sun.com/app/docs/doc/819-0139?q=819-0139>.

Configuring the Sun STMS host settings

Complete this task to configure the Sun Storage Traffic Manager System (STMS) host settings.

To configure the STMS host settings, perform the following steps:

1. Type `cd /etc` to change to the `/etc` subdirectory.
2. Back up the system file in the subdirectory.

3. Edit the system file and set the following parameters for servers with configurations that only use Sun-branded host adapters:

ssd_io_time

This parameter specifies the timeout value for disk operations. Add the following line to the `/etc/system` file to set the `ssd_io_time` parameter for the storage unit LUNs:

```
set ssd:ssd_io_time=0x78
```

ssd_max_throttle

This `ssd_max_throttle` parameter specifies the maximum number of commands that the SSD driver can queue to the host adapter driver. The default value is 256, but you must set the parameter to a value less than or equal to a maximum queue depth for each connected LUN. Determine the value by using the following formula:

$$256 \div (\text{LUNs per adapter})$$

where **LUNs per adapter** is the largest number of LUNs assigned to a single adapter.

To set the `ssd_max_throttle` parameter for the storage unit LUNs in this example, add the following line to the `/etc/system` file:

```
set ssd:ssd_max_throttle=5
```

4. To enable Sun MPxIO multipathing for an STMS host that is attached to a DS6000:
 - a. Type `cd /kernel/drv` to change to the `/kernel/drv` subdirectory.
 - b. Edit the `scsi_vhci.conf` file, and set the parameter values to those in Table 21.

Table 21. Recommended Sun settings for the SCSI device

Parameters	Recommended settings
mpxio-disable	no
load-balance	round-robin
auto-failback	enable

- c. Add the following lines to the end of the `scsi_vhci.conf` file:

```
device-type-scsi-options-list =  
"IBM    1750500", "symmetric-option";  
symmetric-option = 0x1000000;
```

Note: Ensure that there are exactly five spaces between IBM and the storage device ID.

Attaching a Sun host using Sun Cluster

This section describes how to attach a storage unit to a Sun host using Sun Cluster.

See the Sun Web site at <http://www.sun.com> for the most current documentation.

Restriction:

This restriction applies only to a two-node SunCluster 3.x environment with MPxIO that is using SAN version 4.4.8 or lower. When you set up a two-node SunCluster 3.x environment with MPxIO for multipathing, you must use 1:1 zoning. This means that each HBA

must be in a zone that contains only one HBA and only one storage unit target adapter. To ensure that this relationship is maintained, create host nickname definitions using the GUI to enforce port masking. This allows the host WWPN to log in *only* to the desired storage unit target adapter.

This zoning and port masking must be implemented prior to assigning any storage unit LUNs to the SunCluster host and prior to attempting any SunCluster configuration, for example:

```
Host "Sun1" with 2 HBAs: "qlc0" and "qlc1"
Host "Sun2" with 2 HBAs: "qlc0" and "qlc1"
DS6000 with 4 Target adapters:
  B1-C1 (cpsspc400)
  B2-C1 (cpsspc600)
  B3-C1 (cpsspc500)
  B4-C1 (cpsspc700)
```

The following example shows an acceptable zoning configuration:

```
Sun1:qlc0 -- B1-C1
Sun1:qlc1 -- B4-C1
Sun2:qlc0 -- B2-C1
Sun2:qlc1 -- B3-C1
```

Note that this configuration uses only two target adapters, but each is paired with the HBAs in separate zones, which satisfies the requirement that each initiator be in a zone with itself and only 1 target.

The following example shows another acceptable zoning configuration:

```
zone 1: Sun1:qlc0 -- B1-C1
zone 2: Sun1:qlc1 -- B4-C1
zone 3: Sun2:qlc0 -- B2-C1
zone 4: Sun2:qlc1 -- B3-C1
```

This restriction is unique to *two-node* SunCluster 3.x. with MPxIO. The use of MPxIO is limited to Sun-branded host adapters, thus the restriction does not apply to other adapter brands (AMCC, Emulex, QLogic).

IBM will add more information as it becomes available. You can refer to the following references for more information:

- IBM DS6000 Interoperability Matrix, see <http://www.ibm.com/servers/storage/disk/ds6000/>
- PMH 20238,500,624 (2105-800)
- PMH 20794,500,624 / 21374,500,624 (2042 CNT/Inrange Switch)
- CMT Complaint AQ4862

iSCSI Gateway host attachment

This chapter describes how you attach a iSCSI Gateway host system to a DS6000 with fibre-channel adapters.

This chapter contains the following host attachment sections:

- Attachment overview of the iSCSI Gateway host
- Ethernet adapter attachment considerations for the iSCSI Gateway host

- Configuring for storage for the iSCSI Gateway host
- iSCSI Gateway operation through the IP Service Module

Before you attach a host, review the following information:

- the “General requirements for attaching a host system” on page 29
- the DS6000 *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/> for the most current information and details for hosts, operating systems, adapters, and switches
- A list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter at <http://www.ibm.com/servers/storage/support/config/hba/index.wss>

Notes:

1. An iSCSI Gateway host system can be attached with the following adapter cards:
 - For Ethernet using the TCP/IP protocol:
 - Netfinity 1000SX 1GbE adapter (fibre-channel connection)
 - Alacritech 100/1000 1GbE adapter (CAT5 RJ45 connection)
 - Integrated 10/100 Ethernet connections (CAT5 RJ45 connection)
 - NetGear GA600
2. The steps to install and configure adapter cards are examples. Your configuration might be different.

Attachment overview of the iSCSI Gateway host

This section provides an attachment overview of the iSCSI Gateway host.

The iSCSI Gateway router provides access to storage unit fibre-channel storage over an Ethernet (LAN based) network. Host servers attached through Ethernet connections can attach to assigned LUNs in much the same way as fibre-channel attached servers by using host bus adapters. LUNs can shared or dedicated.

A single SN5420 provides a 100 Mbps connection to a storage area network. You can define up to four SCSI routers for the SN5420. You can have up to thirty-two hosts on each SCSI router. Each host can have up to eight storage targets. Each target can present multiple LUNS. For example, you can define eight LUNS for Windows 2000 or Windows 2003. You can define thirty-two LUNs for the UNIX operating system or the Linux operating system. A Windows iSCSI host can discover sixty-four storage devices or eight targets times eight LUNs. A UNIX or Linux host can discover 256 Storage devices or eight targets times thirty-two LUNs.

Note: When you configure the storage unit for host attachment, you can define the worldwide port name by host type. For example, you can define the SN5420 as a iSCSI Gateway host type. You can identify each host by an entry in the scsirouter list. You can present the storage unit LUNs to the SN5420, and you can use the scsirouter list to manage a host assignment.

Ethernet adapter attachment considerations for the iSCSI Gateway host

This section provides Ethernet adapter attachment considerations for the iSCSI Gateway host.

The iSCSI software driver resides on each host (initiator) and completes the connection to a storage device (target). You can only use the iSCSI software drivers (initiators) for the following host systems:

- Windows 2000
- Windows 2003
- Linux (Red Hat Enterprise Linux Server 2.1, or later)

The iSCSI protocol is based on TCP/IP. The iSCSI protocol depends on TCP/IP addressing to identify the SCSI router lists and storage targets. You must install an Ethernet adapter Network Interface Card (NIC) in the host computer.

You can connect NICs of 10/100 or 10/100/1000 through CAT5 copper or fibre connector cables.

Note: For a list of open-systems hosts, the operating systems, adapters and switches that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds6000/>.

Configuring for storage for the iSCSI Gateway host

This section provides information for configuring storage for the iSCSI Gateway host.

You must configure the storage unit to present storage LUNs to the storage router. You must configure the storage router to list the storage LUNs in the SCSI router definition. You must assign a TCP/IP address to the SCSI router. You must identify an entry in the access list table for the storage router for each host. You can group hosts that are listed in the access list table by operating system. The iSCSI software driver must have the TCP/IP address that points to the scsirouter target, and the list of storage LUNs defined in the scsirouter entry.

iSCSI Gateway operation through the IP Service Module

This section discusses the support provided by the IP Service Module.

iSCSI Gateway support is available through the use of the IP Service Module. The IP Service Module supports the following host systems:

- Intel hosts running the Red Hat Enterprise Linux Advanced Server 2.1, or later, operating system
- Intel hosts running the Microsoft® Windows 2000 operating system
- Intel hosts running the Microsoft Windows Server 2003 operating system

Chapter 7. Network configurations

When you are setting up your storage units, you must ensure that you have configured your network so that it can support all the components and functions that you use with your storage unit.

Configuring your network involves ensuring that various components of the storage unit can connect to other points on your local area network (LAN) or outside your LAN. It also involves supporting the ability of other network components, such as hosts, to connect to the storage unit.

Use the following steps as guidelines for configuring your network to support the storage units:

1. Ensure that you meet the network requirements for the Management Consoles.
2. Set up the necessary network connections that support the hosts that will attach to your storage unit.
3. Determine what you need to support the requirements for the modem and for an outside connection if you will use remote support.
4. When your storage units connect to a fibre-channel storage area network (SAN), be sure that you meet the SAN requirements.

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.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the 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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been

estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Accessibility

Accessibility features provide users who have disabilities with the ability to successfully access information and use technology.

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features

These are the major accessibility features in the IBM System Storage DS6000 information:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. IBM Home Page Reader version 3.0 has been tested.
- You can operate features using the keyboard instead of the mouse.

Navigating by keyboard

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the IBM System Storage DS6000 information from the keyboard by using the shortcut keys for your browser or Home Page Reader. See your browser Help for a list of shortcut keys that it supports. See the following Web site for a list of shortcut keys supported by Home Page Reader: http://www-306.ibm.com/able/solution_offerings/keyshort.html

Accessing the publications

You can find HTML versions of the IBM System Storage DS6000 information at the following Web site: <http://www.ehone.ibm.com/public/applications/publications/cgi-bin/pbi.cgi>

You can access the information using IBM Home Page Reader 3.0.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

- AIX
- DB2
- DFSMS/MVS
- DFSMS/VM
- DS4000
- DS6000
- DS8000
- e (logo)
- Enterprise Storage Server
- ES/9000
- ESCON
- FICON
- FlashCopy
- Graphically Dispersed Parallel Sysplex
- HACMP
- i5/OS
- IBM
- IntelliStation
- MVS/ESA
- Netfinity
- NetVista
- Operating System/400
- OS/400
- RS/6000
- S/390
- Seascape
- SNAP/SHOT
- SP
- System/390
- System p5
- System Storage
- Versatile Storage Server
- Virtualization Engine
- VSE/ESA
- z/Architecture
- z/OS
- z/VM
- zSeries

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Other company, product, and service names may be trademarks or service marks of others.

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these Publications for your personal, non commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these Publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING

BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission (FCC) statement

This equipment has been tested and complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the users authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conform à la norme NMB-003 du Canada.

European Union EMC Directive conformance statement

This product is in conformity with the protection requirements of EU Council Directive 2004/108/EC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to CISPR 22/European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

European community contact:

IBM Technical Regulations
Pascalstr. 100, Stuttgart, Germany 70569
Telephone: 0049 (0)711 785 1176
Fax: 0049 (0)711 785 1283
E-mail: tjahn@de.ibm.com

Germany compliance statement

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit.

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 89/336/EWG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der IBM gesteckt/eingebaut werden.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden: "Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 89/336/EWG in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) vom 18. September 1998 (bzw. der EMC EG Richtlinie 89/336) für Geräte der Klasse A.

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Konformitätserklärung nach Paragraph 5 des EMVG ist die IBM Deutschland GmbH, 70548 Stuttgart.

Informationen in Hinsicht EMVG Paragraph 4 Abs. (1) 4:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.

Japanese Voluntary Control Council for Interference (VCCI) class A statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean Ministry of Information and Communication (MIC) statement

Please note that this device has been certified for business use with regard to electromagnetic interference. If you find this is not suitable for your use, you may exchange it for one of residential use.

Taiwan class A compliance statement

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

VS07171L

Index

A

- accessibility 194
- adapter cards
 - installing PCI-FC-1P-OPT 153
 - installing PCI-FC-1P-OPT-A 153
- adapter driver
 - download 30
- adding
 - arrays to ranks 19
 - ranks to extent pools 20
- AdvFS parameters, Tru64 UNIX 65
- AIX
 - host attachment package (pSeries) 75, 76
 - troubleshooting maximum transfer size 78
 - verifying configuration (pSeries) 77
- AlphaServer OpenVMS host
 - requirements for attaching 43
- AMCC adapter
 - Sun host installation 165
- AMCC driver download 30
- AMD host running Linux
 - requirements for attaching 98
- AMD host running VMware
 - requirements for attaching 127
- AMD host running VMware ESX Server
 - installing Emulex adapter 127
- Apple Macintosh host
 - requirements for attaching 35
- applying
 - a configuration 17
- array
 - viewing array DDM information 19
- arrays
 - adding to ranks 19
 - creating 10
 - deleting 26
- attachment, multipath 68
- automatic updates
 - disabling on Linux (i5 server) 71
 - disabling on Linux (pSeries) 81

B

- boot considerations for IBM pSeries hosts, fibre-channel 79
- boot support
 - Windows host 147

C

- cable
 - installing SGI 154
 - SGI installing 154
- channel
 - director 87
- clustering
 - HP-UX host 41

- commands
 - DS CLI
 - mkhostconnect 50
 - Linux
 - fdisk 116
 - mke2fs 118
 - mkfs 118
 - OpenVMS
 - SHOW DEVICES DG 52
 - show system 44
 - pSeries
 - cfgmgr 78
 - importvg 81
 - lscfg -vl fcs 77
 - lsdev -Cc disk 81
 - mount 81
 - smit install_update 75
 - varyonvg 78
 - varyonvg VGname 81
 - SDD
 - datapath query wwpn 77
 - set mode diag 55
 - SGI
 - #/usr/bin/fx -x -d 160
 - SGI host
 - /sbin/scsifo 158
 - sizer -v 55
 - VMware ESX host
 - pbind.pl -A 133
 - vmkfstools 135
 - wwidmgr -show adapter 55
- configuration
 - applying 17
 - network 191
- configuring
 - drivers for SGI 154
 - host device drivers for PRIMEPOWER 38
 - Tru64 UNIX advfs parameters 65
 - Tru64 UNIX kernel parameters 60
 - Configuring and Emulex adapter 102
 - configuring STMS 186
 - configuring storage
 - Tru64 UNIX host 62
 - configuring Tru64 UNIX 60
 - confirming
 - storage connectivity, SGI 156, 157
 - considerations for IBM pSeries hosts, fibre-channel boot 79
- copy services
 - VMware ESX host 136
- Copy Services
 - attachment restrictions 33
- creating
 - arrays 10
 - extent pools 12
 - open systems volume groups 16
 - open systems volumes 14
 - ranks 11
 - zSeries volumes 13
 - creating 13

- creating file systems on Linux 118
- creating LCUs 13

D

- data
 - restoring 81
 - saving 80
- DDM
 - viewing array DDM information 19
- deleting
 - arrays 26
 - extent pools 27
 - host systems 25
 - logical control units (LCUs) 28
 - open systems volumes 27
 - ranks 26
 - volume groups 26
 - zSeries volumes 28
- deleting a configuration 25
- deleting a logical configuration 25
- device drivers, configuring for Sun 167
- director, channel 87
- disks, partitioning
 - Linux 116
- displaying
 - storage unit volume on Tru64 58
- download
 - AMCC adapter driver 30
 - Emulex adapter driver 30
 - HP adapter driver 30
 - JNI adapter driver 30
 - Netfinity adapter driver 30
 - QLogic adapter driver 30
- drivers
 - adapter for Novell NetWare 151
 - configuring SGI 154

E

- Emulex adapter 101, 102
 - Sun host 165
 - VMware ESX Server 127
- Emulex adapter card
 - installing
 - PRIMEPOWER host system 36
- Emulex adapter driver
 - installing 149
- Emulex driver download 30
- eServer BladeCenter HS20 119
- express Configure
 - iSeries volumes 5
 - open systems volumes 3
 - zSeries volumes 4
- extent pools
 - adding ranks 20
 - creating 12
 - deleting 27
 - modifying 21

F

- FAST!Util 119
- FCP for a zSeries host, Accessing the storage unit using 89
- fibre-channel
 - i5/OS Version 5 Release 3 66
 - intermix mode 87
 - iSeries host 66
 - multipath requirement 66
 - Novell NetWare host 149
 - pSeries host 74
 - requirement, multipath 66
 - RS/6000 host 74
 - S/390 85
 - Version 5 Release 3, i5/OS 66
 - Windows 2000 host 137
 - Windows 2003 host 137
 - zSeries 85
- fibre-channel arbitrated loop, SGI 157
- fibre-channel boot
 - Intelfibre-channel boot
 - Linux 119
 - Linux
 - Intel 119
- FICON
 - S/390 85
 - S/390 host 87
 - S/390 host, overview 86
 - zSeries 85
 - zSeries host 87
 - zSeries host (overview) 86
- file systems on Linux 118
- FlashCopy
 - attachment restrictions 33
- Fujitsu PRIMEPOWER host
 - fibre-channel 35
 - WWPN 36

G

- Gateway 500 host
 - requirements for attaching 95
- guest partition
 - Linux 72

H

- HACMP
 - multiple (pSeries) 80
 - pSeries 80
- Hewlett-Packard host
 - locating WWPN 40
- Hewlett-Packard Tru64 UNIX host
 - WWPN 57
- host attachment
 - general requirements 29
- host attachment package
 - pSeries or RS/6000 hosts (AIX) 75, 76
- host systems
 - deleting 25
- hosts, fibre-channel boot considerations
 - for IBM pSeries 79
- hosts, Registered state-change notifications (RSCNs) on zSeries 94

- HP AlphaServer
 - KGPSAxx/FCA2xxx adapter 44
- HP AlphaServer OpenVMS host
 - requirements for attaching 43
- HP driver download 30
- HP OpenVMS host
 - defining UDID 49
 - WWPN 45
- HP Tru64 UNIX host
 - requirements for attaching 54
- HP-UX host
 - MC/ServiceGuard 41
 - queue depth 41

I

- i5/OS Version 5 Release 3 68
- IBM NAS Gateway 500
 - WWPN 96
- IBM NAS Gateway 500 host
 - requirements for attaching 95
- IBM Subsystem Device Driver, Installing Linux over the SAN with an 121
- installing
 - adapter driver
 - Novell NetWare 151
 - device driver for a Netfinity adapter card
 - Windows 143
 - Emulex adapter card 36
 - Emulex adapter driver 149
 - fibre-channel adapter drivers for SGI 154
 - optical cable, SGI 154
 - optical SGI cable 154
 - QLogic adapter card 166
 - QLogic QLA23xx adapter card 143
 - SGI PCI-FC-1P-OPT 153
 - SGI PCI-FC-1P-OPT-A 153
- installing Emulex adapter 101
- installing on Linux Intel host 103, 105
- installing on Linux Intel or AMD host 101
- installing QLogic adapter 103, 105
- Integrity server OpenVMS host
 - requirements for attaching 43
- Intel host running Linux
 - requirements for attaching 98
- Intel host running VMware
 - requirements for attaching 127
- Intel host running VMware ESX Server
 - installing Emulex adapter 127
- Interoperability Matrix 29
- IRIX multipathing
 - SGI 158
- IRIX, SGI host
 - checking OS version 153
- iSCSI Gateway host
 - attachment considerations 190
 - overview 189
 - requirements for attaching 188
- iSeries
 - WWPN 67
- iSeries host
 - configuration 68
 - feature code, configuration by 68
 - fibre-channel 66

- iSeries host (*continued*)
 - LUN considerations 67
 - switches 68
- iSeries volumes
 - creating 5

J

- JNI driver download 30

K

- keyboards
 - accessibility features 194
- KGPSA-xx/FCA2xxx adapter
 - installing on Tru64 55
- KGPSAxx/FCA2xxx adapter
 - installing on HP AlphaServer 44
 - setting mode on HP AlphaServer 45

L

- LCUs
 - creating 13
 - modifying 23
- legal
 - terms and conditions 196
- lilo 121
- limitations for Tru64 UNIX 65
- Linux
 - disable automatic updates (i5 server) 71
 - disable automatic updates (pSeries) 81
 - file systems 118
 - guest partition 72
 - installing over a SAN without SDD 119
 - SDD
 - installing on Linux (pSeries) 81
 - Linux host 106
 - Linux Intel host 103, 105
 - Linux Intel or AMD host 101

M

- MC/ServiceGuard
 - HP-UX, configuration 41
- mke2fs 118
- mkfs 118
- modifying
 - extent pools 21
 - LCUs 23
 - open systems volume groups 21
 - zSeries volumes 22
- modifying ranks 20
- multipath attachment 68

N

- NAS Gateway 500 host
 - WWPN 96
- NAS Gateway 500, WWPNs using a Web browser 96

- Netfinity
 - Windows 137
- Netfinity adapter
 - Windows 141
- Netfinity adapter card
 - installing Windows device driver 143
- Netfinity driver download 30
- network, configuration 191
- Novell Netware host
 - QLogic adapter, installin 150
- Novell NetWare host
 - installing drivers 151
 - requirements for attaching 149
 - WWPN 152

O

- open system volumes
 - deleting 27
- open systems volume groups
 - creating 16
- open systems volumes
 - creating 3, 14
- OpenVMS
 - switch connectivity 48
- OpenVMS host
 - accessing volumes 51
 - confirming operating system installation 43
 - manually configure volumes 51
- optical cable, installing SGI 154

P

- parameters
 - recommended configuration for AMCC adapters on a Sun host 170
 - recommended configuration for Emulex adapters on a Sun host 172
 - recommended configuration for QLogic adapters on a Sun host 177, 178
 - recommended configuration for Sun-Emulex adapters on a Sun host 174
 - recommended configuration for Sun-QLogic adapters on a Sun host 180
 - settings for Emulex LP9002L 36
- partitioning storage unit disks
 - Linux 116
- PCI-FC-1P-OPT adapter card, installing SGI 153
- persistent binding
 - VMware ESX host 133
- persistent reserves, removing Tru64 UNIX 62
- primary path, switching I/O 160
- properties
 - viewing and modifying 19
- pSeries
 - SAN changes 78
- pSeries host 79
 - attaching multiple (without HACMP) 80
 - attaching multiple with HACMP 80

- pSeries host (*continued*)
 - disabling automatic updates 81
 - fibre-channel boot 79
 - host attachment package (AIX) 75, 76
 - requirements for attaching 74
 - running Linux 81
 - installing SDD 81
 - verifying configuration 77
 - WWPN 77
- pSeries hosts, fibre-channel boot considerations for IBM 79
- public and private loop modes
 - Emulex for a PrimePower host 36

Q

- qfull_retry_count 102
- QLA23xx adapter card, installing 143
- QLogic
 - adapter driver 121
 - Emulex
 - Windows 137
 - Windows 137
- QLogic adapter 103, 105, 119
- VMware ESX Server 128
- QLogic adapter card, installing 150, 166
- QLogic driver download 30
- QLogic QLA23xx adapter card, installing 143
- queue depth
 - HP-UX host 41

R

- rank
 - repairing 20
- ranks
 - adding arrays to 19
 - creating 11
 - modifying 20
- registry, setting TimeOutValue for Windows 146
- Remote Mirror and Copy
 - attachment restrictions 33
- removing a rank from an extent pool 21
- removing persistent reserves, Tru64 UNIX 62
- requirements for attaching 137
 - Fujitsu PRIMEPOWER host 35
 - HP-UX host 40
 - iSeries host 66
 - S/390 host 85
 - Sun host 164
 - zSeries host 85
- restoring data on storage unit 81
- RS/6000 host
 - attaching multiple (without HACMP) 80
 - attaching multiple with HACMP 80
 - host attachment package (AIX) 75, 76
 - requirements for attaching 74
 - verifying configuration 77
 - WWPN 77

S

- S/390 host
 - FICON 86
 - FICON channel attachment 87
 - service information messages 95
- SAN changes
 - pSeries 78
- SAN File System host
 - requirements for attaching 97
- SAN Volume Controller host
 - requirements for attaching 98
- saving data on the storage unit 80
- SCSI parameters, Tru64 UNIX 60
- secondary path, switching I/O 160
- server
 - restoring data on the storage unit 81
 - saving data on the storage unit 80
- service information messages
 - zSeries and S/390 host 95
- settings
 - parameters for Emulex LP9002L 36
 - recommended configuration for AMCC adapters on a Sun host 170
 - recommended configuration for Emulex adapters on a Sun host 172
 - recommended configuration for QLogic adapters on a Sun host 178
 - recommended configuration for Sun-Emulex adapters on a Sun host 174
 - recommended configuration for Sun-QLogic adapters on a Sun host 180
- SGI
 - commands
 - #/usr/bin/fx -x -d 160
 - fibre-channel arbitrated loop 157
 - IRIX multipathing 158
 - switched fabric 156
 - WWPN 156
- SGI host
 - checking IRIX version 153
 - configuring for host failover 158
 - configuring storage 160, 162
 - confirming availability of failover 158
 - confirming connectivity 156
 - confirming switch connectivity 154
 - displaying zoning information 155
 - fibre-channel arbitrated loop 160
 - see Silicon Graphics host 152
 - switched fabric 159
 - verifying installation of adapter card 153
- Silicon Graphics host
 - requirements for attaching 152
- state-change notifications (RSCNs) on zSeries hosts, Registered 94
- status
 - viewing 19
- storage connectivity
 - Tru64 UNIX 58
- storage plex
 - configuring iSeries volumes quickly 5
 - configuring open systems volumes quickly 3

- storage plex (*continued*)
 - configuring zSeries volumes quickly 4
- storage server
 - restoring data 81
 - saving data 80
- storage unit
 - restoring data on the 81
 - saving data 80
- storage unit volume
 - storage unit volume on Tru64 58
- Subsystem Device Driver, Installing Linux over the SAN 121
- Sun adapter 187
 - Sun host 166
- Sun host 186
 - AMCC adapter 165
 - configuring host device drivers 167
 - Emulex adapter 165
 - recommended configuration for
 - QLogic adapters on a Sun host 177
 - setting system parameters for AMCC adapters 183
 - setting system parameters for Emulex adapters 184
 - setting system parameters for QLogic adapters 184
 - setting system parameters for
 - Sun-Emulex adapters 185
 - setting system parameters for
 - Sun-QLogic adapters 185
 - Sun adapter 166
 - WWPN 185
- Sun-branded adapter 187
- switch connectivity
 - OpenVMS 48
 - Tru64 UNIX 57
- switched fabric, SGI 156
- switches
 - iSeries host 68
- switching I/O between the primary and secondary alternate path, SGI 160

T

- TimeOutValue registry setting for
 - Windows 146
- Trademarks 195
- troubleshooting 102
 - qfull_retry_count 105
 - QLogic adapter 105
 - service information messages 95
- Tru64
 - installing KGPSA-xx/FCA2xxx adapter 55
- Tru64 UNIX
 - limitations 65
 - storage connectivity 58
 - switch connectivity 57
- Tru64 UNIX host
 - configuring storage 62
 - confirming OS install 55
 - verifying attachment 61

U

- UDID, HP OpenVMS 49
- updating
 - Windows device driver 145

V

- verifying
 - host configured for storage 146, 190
- Version 5 Release 3, i5/OS 68
- viewing and modifying properties 19
- VM file system
 - VMware ESX host 135
- VMware ESX host
 - commands
 - pbind.pl -A 133
 - copy services considerations 136
 - VM file system 135
- VMware ESX Server
 - Emulex adapter 127
 - QLogic adapter 128
- volume groups
 - deleting 26

W

- Windows
 - Netfinity adapter 141
- Windows 2000 host 137
- Windows 2003 host 137
- Windows device driver, updating
 - the 145
- Windows host
 - boot support 147
 - WWPN 145
- WWPN 106
 - Hewlett-Packard host 40
 - Hewlett-Packard Tru64 UNIX host 57
 - HP OpenVMS host 45
 - iSeries host 67
 - Linux host 106
 - NAS Gateway 500 host 96
 - Novell NetWare host 152
 - pSeries host 77
 - RS/6000 host 77
 - SGI host 156
 - Sun host 185
 - Windows host 145
- WWPNs, NAS Gateway 500 96

Z

- zoning
 - iSeries switch 68
- zSeries
 - volumes
 - deleting 28
- zSeries host
 - FICON 86
 - FICON channel attachment 87
 - service information messages 95
- zSeries host, Accessing the storage unit
 - using FCP for a 89
- zSeries hosts, Registered state-change notifications (RSCNs) on 94

- zSeries volumes
 - creating quickly 4
 - modifying 22



Printed in USA