IBM System Storage DS8000
Version 6 Release 2

Host Systems Attachment Guide

GC27-2298-02

IBM System Storage DS8000
Version 6 Release 2

Host Systems Attachment Guide

GC27-2298-02

Note:
Before using this information and the product it supports, read the information in the **Safety and environmental notices** and **Notices** sections.

This edition applies to version 6, release 2 of the IBM System Storage Host Attachment Guide, and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Figures

Tables

About this guide

This guide provides information about attaching hosts to your DS8000® storage unit. The first chapter provides an overview of host attachment concepts, specific considerations that affect more than one host. Each subsequent chapter provides information on a specific host. Table 1 lists the hosts that you can attach to your storage unit and provides a link to the chapter for each host.

For the most current information on supported hosts, operating systems, adapters and switches, go to the *System Storage® Interoperation Center* (SSIC) website at IBM Storage SSIC.

Note: SAN Volume Controller information is available on the SAN Volume Controller interoperability website (http://www.ibm.com/systems/storage/software/virtualization/svc/interop.html).

Table 1. Hosts that are supported by the DS8000

Host	Chapter
Apple Macintosh	Chapter 2, "Apple Macintosh host attachment," on page 15
Fujitsu PrimePower	Chapter 3, "Fujitsu PrimePower host attachment," on page 19
HP OpenVMS	Chapter 4, "HP OpenVMS host attachment," on page 21
HP Tru64 UNIX	Chapter 5, "HP Tru64 UNIX host attachment," on page 33
HP-UX	Chapter 6, "HP-UX host attachment," on page 47
IBM® Power Systems™ (IBM i, AIX, Linux)	Chapter 7, "IBM Power Systems host attachment," on page 51
IBM SAN Volume Controller	Chapter 8, "IBM SAN Volume Controller host attachment," on page 69
IBM NAS Gateway 500	Chapter 9, "IBM NAS Gateway 500 host attachment," on page 71
IBM System z® or IBM S/390®	Chapter 10, "IBM System z or IBM S/390 host attachment," on page 73
Linux (x86 or x86_64)	Chapter 11, "Linux (x86 or x86_64) host attachment," on page 95
Microsoft Windows Server	Chapter 12, "Microsoft Windows Server host attachment," on page 115
VMware ESX Server	Chapter 13, "VMware ESX Server host attachment," on page 121
Novell NetWare	Chapter 14, "Novell NetWare host attachment," on page 131
Silicon Graphics (SGI)	Chapter 15, "Silicon Graphics host attachment," on page 137
Sun	Chapter 16, "Sun host attachment," on page 151

Conventions used in this guide

The following typefaces are used to show emphasis:

boldface
Text in **boldface** represents menu items and lowercase or mixed-case command names.

italics Text in *italics* is used to emphasize a word. In command syntax, it is used for variables for which you supply actual values.

`monospace`
Text in `monospace` identifies the data or commands that you type, samples of command output, or examples of program code or messages from the system.

DS8000 library and related publications

Product manuals, other IBM publications, and websites contain information that relates to DS8000.

DS8000 Information Center

The IBM System Storage DS8000 Information Center contains all of the information that is required to install, configure, and manage the DS8000. The information center is updated between DS8000 product releases to provide the most current documentation. The information center is available at the following website: DS8000 Information Center

DS8000 library

Table 2 lists and describes the publications that make up the DS8000 library. Unless otherwise noted, these publications are available in Adobe portable document format (PDF). Go to the DS8000 Library to obtain a publication.

Table 2. DS8000 library

Title	Description	Order Number
IBM System Storage DS: Command-Line Interface User's Guide	This guide describes the commands that you can use from the command-line interface (CLI) for managing your DS8000 configuration and Copy Services relationships. The CLI provides a set of commands that you can use to write customized scripts for a host system.	GC53-1127
IBM System Storage DS8000: Host Systems Attachment Guide	This guide provides information about attaching hosts to the DS8000 storage unit. The DS8000 provides a variety of host attachments so that you can consolidate storage capacity and workloads for open-systems hosts and System z or S/390 hosts.	GC27-2298
IBM System Storage DS8000: Introduction and Planning Guide	This guide introduces the DS8000 product and lists the features you can order. It also provides guidelines for planning the installation and configuration of the storage unit.	GC27-2297

Table 2. DS8000 library (continued)

Title	Description	Order Number
IBM System Storage Multipath Subsystem Device Driver User's Guide	This publication describes how to use the IBM Subsystem Device Driver (SDD) on open-systems hosts to enhance performance and availability on the DS8000. SDD creates single devices that consolidate redundant paths for logical unit numbers. SDD permits applications to run without interruption when path errors occur. It balances the workload across paths, and it transparently integrates with applications.	GC27-2122
IBM System Storage DS Application Programming Interface Reference	This publication provides reference information for the IBM System Storage DS application programming interface (API) and provides instructions for installing the Common Information Model Agent, which implements the API.	GC35-0516

Other IBM publications

Other IBM publications contain additional information that is related to the DS8000 product library. Table 3 is divided into categories to help you find publications that are related to specific topics.

Table 3. Other IBM publications

Title	Description	Order number
System Storage Productivity Center		
IBM System Storage Productivity Center Introduction and Planning Guide	This publication introduces the IBM System Storage Productivity Center hardware and software.	SC23-8824
Read This First: Installing the IBM System Storage Productivity Center	This publication provides quick instructions for installing the IBM System Storage Productivity Center hardware.	GI11-8938
IBM System Storage Productivity Center Software Installation and User's Guide	This publication describes how to install and use the IBM System Storage Productivity Center software.	SC23-8823
IBM System Storage Productivity Center User's Guide	This publication describes how to use the IBM System Storage Productivity Center to manage the DS8000, IBM System Storage SAN Volume Controller clusters, and other components of your data storage infrastructure from a single interface.	SC27–2336
IBM Tivoli® Key Lifecycle Manager		
IBM Tivoli Key Lifecycle Manager Installation and Configuration Manager	This publication describes how to install and configure the Tivoli encryption key manager. The key server can be used to manage the encryption keys assigned to the IBM Full Disk Encryption disk drives in the DS8000.	SC23-9977
IBM System Management Pack for Microsoft		

Table 3. Other IBM publications (continued)

Title	Description	Order number
IBM System Management Pack for Microsoft System Center Operations Manager User Guide	This publication describes how to install, configure, and use the IBM Storage Management Pack for Microsoft System Center Operations Manager (SCOM).	GC27-3909

IBM documentation and related websites

The following websites provide information about the DS8000 or related products or technologies:

Table 4. IBM documentation and related websites

Website	Link
IBM System Storage DS8000 series	IBM Storage Disk DS8000
Support for DS8000, IBM System Storage, and IBM TotalStorage® products	IBM TotalStorage Products Support
Concurrent Copy for IBM System z and S/390 host systems	IBM Storage Software System z and S/390 host systems
DS8000 command-line interface (DS CLI)	DS8000 Information Center The information center has a complete command reference for the DS CLI.
Information about code bundles for DS8700 and DS8800.	IBM Support - DS8700 code bundles See **Section 3** for cross-reference links to SDD. IBM Support - DS8800 code bundles
IBM FlashCopy® for System z and S/390 host systems	IBM Storage Software System z and S/390 host systems
Host system models, operating systems, adapters, and switches that the DS8000 series supports	Host system models, OS, adapters and switches Click **Interoperability matrix**. IBM Storage SSIC Click **New search**.
IBM Disk Storage Feature Activation (DSFA)	IBM Storage DSFA
IBM version of the Java SE Runtime Environment (JRE) that is often required for IBM products	IBM Java SE Runtime Environment (JRE)
Information about IBM Storage Easy Tier™	• IBM Storage Easy Tier Deployment Considerations Guide • IBM Storage Easy Tier performance on DS8700
Remote Mirror and Copy (formerly Peer-to-Peer Remote Copy [PPRC]) for System z and S/390 host systems	IBM Storage Software System z and S/390 host systems
SAN Fibre Channel switches	IBM Storage SAN switches
Subsystem Device Driver (SDD)	IBM Storage Software SDD

Table 4. IBM documentation and related websites (continued)

Website	Link
IBM Publications Center	IBM Publications Center
IBM Redbooks® publications	IBM Redbooks Publications

Related accessibility information

To view a PDF file, you need Adobe Acrobat Reader, which can be downloaded for free from the Adobe website at: Adobe Downloads

How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

To submit any comments about this book or any other DS8000 documentation:

- Go to the DS8000 information center website at DS8000 Feedback. There you will find the feedback page where you can enter and submit comments.
- Send your comments by email to starpubs@us.ibm.com. Be sure to include the following information:
  - In the subject line of the email:
    - Exact publication title and version
    - Publication form number (for example, GC26-1234-02)
  - Page, table, or illustration numbers that you are commenting on
  - A detailed description of any information that should be changed

Safety and Environmental notices

This section contains information about safety notices that are used in this guide and environmental notices for this product.

Safety notices

Observe the safety notices when using this product. These safety notices contain danger and caution notices. These notices are sometimes accompanied by symbols that represent the severity of the safety condition.

Most danger or caution notices contain a reference number (Dxxx or Cxxx). Use the reference number to check the translation in the *IBM System Storage DS8000 Safety Notices*, P/N 98Y1543.

The sections that follow define each type of safety notice and give examples.

Danger notice

A danger notice calls attention to a situation that is potentially lethal or extremely hazardous to people. A lightning bolt symbol always accompanies a danger notice to represent a dangerous electrical condition. A sample danger notice follows:

DANGER: An electrical outlet that is not correctly wired could place hazardous voltage on metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. *(D004)*

Caution notice

A caution notice calls attention to a situation that is potentially hazardous to people because of some existing condition, or to a potentially dangerous situation that might develop because of some unsafe practice. A caution notice can be accompanied by one of several symbols:

If the symbol is...	It means...
	A generally hazardous condition not represented by other safety symbols.
 Class II	This product contains a Class II laser. Do not stare into the beam. (*C029*) Laser symbols are always accompanied by the classification of the laser as defined by the U. S. Department of Health and Human Services (for example, Class I, Class II, and so forth).
	A hazardous condition due to mechanical movement in or around the product.

If the symbol is...	It means...
	This part or unit is heavy but has a weight smaller than 18 kg (39.7 lb). Use care when lifting, removing, or installing this part or unit. (*C008*)

Sample caution notices follow:

Caution

The battery is a lithium ion battery. To avoid possible explosion, do not burn. Exchange only with the IBM-approved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (*C007*)

Caution

The system contains circuit cards, assemblies, or both that contain lead solder. To avoid the release of lead (Pb) into the environment, do not burn. Discard the circuit card as instructed by local regulations. (*C014*)

Caution

When removing the Modular Refrigeration Unit (MRU), immediately remove any oil residue from the MRU support shelf, floor, and any other area to prevent injuries because of slips or falls. Do not use refrigerant lines or connectors to lift, move, or remove the MRU. Use handholds as instructed by service procedures. (*C016*)

Caution

Do not connect an IBM control unit directly to a public optical network. The customer must use an additional connectivity device between an IBM control unit optical adapter (that is, fibre, ESCON®, FICON®) and an external public network . Use a device such as a patch panel, a router, or a switch. You do not need an additional connectivity device for optical fibre connectivity that does not pass through a public network.

Environmental notices

The environmental notices that apply to this product are provided in the *Environmental Notices and User Guide*, Z125-5823-*xx* manual. A copy of this manual is located on the publications CD.

Summary of changes

This document contains terminology, maintenance, and editorial changes for version GC27-2298-02 of the IBM System Storage DS8000 Host Systems Attachment Guide. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

New information

- Added procedure for configuring Red Hat Enterprise Linux (RHEL) 6 or later.

Changed Information

- Updated the following chapter names:
  - Chapter 11: *Intel or AMD running Linux host attachment* is now *Linux (x86 or x86_64) host attachment*.
  - Chapter 12: *Intel or AMD running Microsoft Windows Server host attachment* is now *Microsoft Windows Server host attachment*.
  - Chapter 13: *Intel or AMD running VMware ESX Server host attachment* is now *VMware ESX Server host attachment*.
  - Chapter 14: *Intel running Novell NetWare host attachment* is now *Novell NetWare host attachment*.
- Updated table of supported HBAs for Linux on System z to include FICON Express4 and FICON Express8S on System z9®, System z10™, System z114.
- Updated introductory information at the beginning of Chapter 11.
- Updated information on dynamic SAN fabric reconfiguration and renamed the topic *Rescanning the SAN fabric*. Deleted references to unloading and reloading the host adapter driver and manually adding and removing SCSI disks by echoing the /proc or /sys filesystem.
- Updated parameter settings for QLogic QLA23*xx* adapters on a Sun host to support driver level v5.05 build 1.

Deleted information

- Removed sections on disabling automatic system updates for Linux on IBM i and running Linux in a guest partition on an IBM i host.
- Removed sections on installing Emulex and QLogic adapters on an Intel or AMD host running Linux.
- Removed section on configuring and troubleshooting an Emulex adapter on an Intel or AMD host running Linux.
- Removed sections on installing and updating Emulex and QLogic drivers on an Intel or AMD host running Linux.
- Removed procedures for creating a module disk for SUSE Linux Enterprise Server 9.0, installing Linux over the SAN with and without SDD, and updating a module without SDD.
- Removed SCSI disk limitations on Linux.
- Removed references to Microsoft Windows Server 2000, which is no longer supported.
- Removed references to Emulex LP9402DC adapter, which is no longer supported.

- Removed recommended configuration file parameters and StorPort Miniport driver settings for Emulex adapters on Microsoft Windows hosts, because the system uses default parameters.
- Removed recommended settings for a QLogic adapter on a VMware ESX Server host because the system uses default parameters.

Chapter 1. Introduction

This chapter is an overview of host attachment for the DS8000 series.

This chapter contains the following sections:

General requirements for attaching a host

Before you attach a host to the DS8000, review this list of general requirements for all hosts. The remaining chapters of this guide detail specific host requirements for each host type.

Perform the following steps before you attach any host system to a DS8000:

1. Go to the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC for the most current information on supported hosts, operating systems, adapters, and switches.
2. Obtain a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter on the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.
3. Ensure that you can reference the following documentation:
   - The IBM System Storage DS8000 Information Center at DS8000 Information Center.
   - The IBM Systems Information Centers at publib.boulder.ibm.com/eserver/.
4. Have an IBM service support representative install the storage unit.
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" on page 11.
6. Review the "I/O adapter features" section of the *IBM System Storage DS8000 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 web-based 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 "Host configuration using the DS8000 interfaces" on page 13 for more information.
8. Install the adapter driver with your host adapter or use the steps defined in "Downloading and installing a host bus adapter driver" on page 2 to download and install an updated adapter driver.
9. For further host attachment details, refer to the subsequent chapter dedicated to the host OS.

Downloading and installing a host bus adapter driver

Complete this task to download and install host adapter drivers to use with the DS8000.

About this task

To download and install a host adapter driver, perform the following steps:

Notes:

1. For Linux hosts, use the HBA driver provided in the distribution.
2. You must download the relevant vendor documentation for the driver that you select to correctly install and configure the host adapter.

Procedure

1. Go to the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC and verify that the host adapter you want to use is compatible with the DS8000 series, your host, and your host operating system.
2. Obtain a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter on the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.
3. If required, download the adapter driver from the adapter vendor website. Table 5 provides vendor web addresses and specific download steps. For the most current steps, see the vendor website download page.

Table 5. Host adapter driver download steps by vendor

Host adapter vendor	Web address	Steps to locate download page
AMCC/JNI	www.amcc.com	1. Click **Downloads → FC HBA/OEM Files**. 2. Click **IBM**. 3. Locate your adapter. 4. Download the appropriate files for your adapter.
Emulex	www.emulex.com/downloads/ibm/oneconnect-software-kits.html	1. Click the tab for the adapter type that you are looking for. 2. Locate your adapter. 3. Download the appropriate files for your adapter.
Hewlett-Packard	www.hp.com	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 **Download** to download the adapter driver. 4. Return to the search results and review the links to documentation, such as installation requirements and release notes.
QLogic	support.qlogic.com/support/oem_ibm.asp	1. Click the DS8000 link. 2. Locate your adapter. 3. Download the appropriate files for your adapter.

Table 5. Host adapter driver download steps by vendor (continued)

Host adapter vendor	Web address	Steps to locate download page
Sun	www.sun.com/storage/san	If you are using Solaris 8 or 9, perform the following steps. If you are using Solaris 10, you can skip these steps because the SAN software is included in the operating system. 1. Scroll down to the **Get the Software** section. 2. Locate the current driver in the list and click the appropriate link. 3. Type your **Username** and **Password** in the fields, and click **Log in and Continue**. If you do not have a user name and password, complete the registration process and return to this page when you have received them. **Note:** Sun only grants user names and passwords to customers who have purchased maintenance contracts. 4. Click **Accept** to agree to the license agreement (required). 5. Download the appropriate files for your adapter.

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 web address that is listed in Table 5 on page 2 to locate installation and configuration instructions.
5. Locate the chapter for your host in the *IBM System Storage DS8000 Host Systems Attachment Guide*, and follow any additional driver configuration instructions.

Multipathing

Installing a multipathing driver on the host is highly recommended as it provides redundancy for failed paths and failed storage controllers. The DS8000 supports a variety of multipathing drivers, depending upon the OS. Please refer to the System Storage Interoperation Center (SSIC) website at www.ibm.com/systems/support/storage/config/ssic/ for the correct multipathing driver.

IBM System Storage Multipath Subsystem Device Driver

The IBM System Storage Multipath Subsystem Device Driver (SDD) is IBM's multipathing driver and is used with several operating systems.

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
- Multiple path SAN boot on storage devices that support multipath I/O (MPIO)

The term *SDD* represents both the SDD as the established multipath subsystem device driver and as the MPIO path control module, depending upon the operating system. Table 6 provides examples.

Table 6. Examples of SDD MPIO path control modules

Operating system	Description	Acronym
AIX®	SDD Path Control Module	SDDPCM
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 IBM Storage Software SDD.

Host attachment path considerations for a storage image

Path considerations affect the performance and availability of a storage image.

For optimal performance, use the following guidelines if you are attaching multiple paths from a single host system to I/O ports on a host adapter of a storage image:

- Use attached I/O ports on different host adapters.
- Use multiple physical adapters.
- Do not use all the ports on each host adapter.

Figure 1 on page 5 is a graphical representation of Rack 1 and Rack 2.

This side belongs to Cluster 0 (Top Cluster)

	Enclosure 0 (cpssbay00)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B00**	B01	1002	B03	B04	1005
Pt	Logical name					
	cpssfc000	cpssfc001	iss002	cpssfc003	cpssfc004	iss005
0	cpssfc0000	cpssfc0010		cpssfc0030	cpssfc0040	
1	cpssfc0001	cpssfc0011		cpssfc0031	cpssfc0041	
2	cpssfc0002	cpssfc0012		cpssfc0032	cpssfc0042	
3	cpssfc0003	cpssfc0013		cpssfc0033	cpssfc0043	

	Enclosure 2 (cpssbay02)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B0C	B0D	100E	B0F	B10	1011
Pt	Logical name					
	cpssfc020	cpssfc021	iss022	cpssfc023	cpssfc024	iss025
0	cpssfc0200	cpssfc0210		cpssfc0230	cpssfc0240	
1	cpssfc0201	cpssfc0211		cpssfc0231	cpssfc0241	
2	cpssfc0202	cpssfc0212		cpssfc0232	cpssfc0242	
3	cpssfc0203	cpssfc0213		cpssfc0233	cpssfc0243	

	Enclosure 4 (cpssbay04)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B18	B19	101A	B1B	B1C	101D
Pt	Logical name					
	cpssfc040	cpssfc041	iss042	cpssfc043	cpssfc044	iss045
0	cpssfc0400	cpssfc0410		cpssfc0430	cpssfc0440	
1	cpssfc0401	cpssfc0411		cpssfc0431	cpssfc0441	
2	cpssfc0402	cpssfc0412		cpssfc0432	cpssfc0442	
3	cpssfc0403	cpssfc0413		cpssfc0433	cpssfc0443	

	Enclosure 6 (cpssbay06)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B24	B25	1026	B27	B28	1029
Pt	Logical name					
	cpssfc060	cpssfc061	**iss062**	cpssfc063	cpssfc064	**iss065**
0	cpssfc0600	cpssfc0610		cpssfc0630	cpssfc0640	
1	cpssfc0601	cpssfc0611		cpssfc0631	cpssfc0641	
2	cpssfc0602	cpssfc0612		cpssfc0632	cpssfc0642	
3	cpssfc0603	cpssfc0613		cpssfc0633	cpssfc0643	

This side belongs to Cluster 1 (Bottom Cluster)

	Enclosure 1 (cpssbay01)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B06**	B07	1008	B09	B0A	100B
Pt	Logical name					
	cpssfc010	cpssfc011	iss012	cpssfc013	cpssfc014	iss015
0	cpssfc0100	cpssfc0110		cpssfc0130	cpssfc0140	
1	cpssfc0101	cpssfc0111		cpssfc0131	cpssfc0141	
2	cpssfc0102	cpssfc0112		cpssfc0132	cpssfc0142	
3	cpssfc0103	cpssfc0113		cpssfc0133	cpssfc0143	

	Enclosure 3 (cpssbay03)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B12	B13	B14	B15	B16	1017
Pt	Logical name					
	cpssfc030	cpssfc031	iss032	cpssfc033	cpssfc034	iss035
0	cpssfc0300	cpssfc0310		cpssfc0330	cpssfc0340	
1	cpssfc0301	cpssfc0311		cpssfc0331	cpssfc0341	
2	cpssfc0302	cpssfc0312		cpssfc0332	cpssfc0342	
3	cpssfc0303	cpssfc0313		cpssfc0333	cpssfc0343	

	Enclosure 5 (cpssbay05)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B1E	B1F	1020	B21	B22	1023
Pt	Logical name					
	cpssfc050	cpssfc051	iss052	cpssfc053	cpssfc054	iss055
0	cpssfc0500	cpssfc0510		cpssfc0530	cpssfc0540	
1	cpssfc0501	cpssfc0511		cpssfc0531	cpssfc0541	
2	cpssfc0502	cpssfc0512		cpssfc0532	cpssfc0542	
3	cpssfc0503	cpssfc0513		cpssfc0533	cpssfc0543	

	Enclosure 7 (cpssbay07)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B2A	B2B	102C	B2D	B2E	102F
Pt	Logical name					
	cpssfc070	cpssfc071	iss072	cpssfc073	cpssfc074	iss075
0	cpssfc0700	cpssfc0710		cpssfc0730	cpssfc0740	
1	cpssfc0701	cpssfc0711		cpssfc0731	cpssfc0741	
2	cpssfc0702	cpssfc0712		cpssfc0732	cpssfc0742	
3	cpssfc0703	cpssfc0713		cpssfc0733	cpssfc0743	

For Fibre Channel Adapters:
- Logical names are cpssfcEES (Enclosure, Slot)
- Port logical names are cpssfcEESP (Enclosure, Slot, Port)
- SAID tags are 0xEESP (Enclosure, Slot, Port)

For ESCON Adapters:
- Logical names are “cpssnpEES“ (Enclosure, Slot)
- Port logical names are “cpssnpEESP“ (Enclosure, Slot, Port)
- Adapter IDs are 8XX instead of BXX.

For Device Adapters:
- logical names are “issEES“ (Enclosure, Slot)
- adapter IDs are 10XX.

FC Adapter ID = 0B00 + EE*6 + slot

** **1750** has 2 HAs: B00/cpssfc000 and B06/cpssfc010
and 2 DAs: 1001/iss001 and 1007/iss011

Enclosure numbers (EE) for expansion frames (from the back):

Rack 5		Rack 4		Rack 3		Rack 2		Rack 1	
10 B60	11 B66	C B48	D B4E	8 B30	9 B36	4 B18	5 B1E	0 B00	1 B06
12 B6C	13 B72	E B54	F B5A	A B3C	B B42	6 B24	7 B2A	2 B0C	3 B12

f2c00943

Figure 1. 2107 Adapter Identification Map. In the figure, Rack 1 includes enclosures 0, 2, 4, and 6. Rack 2 includes enclosures 1, 3, 5, and 7. Each enclosure includes slot numbers 0, 1, 2, 3, 4, and 5 which run along the top, and port numbers 0, 1, 2, and 3 which run down the left side. Slot numbers 2 and 5 are reserved for host attachment. An adapter ID and logical name are associated with each slot number. The adapter ID numbers B00, B01, B03, and B04 are internal microcode hexadecimal numbers. The strings like cpssfc0010, cpssfc0011, cpssfc0020 and cpssfc0042 are logical adapter interface names for the ports on the adapters.

Note: A maximum of two host adapter cards per I/O bay can be used on a DS8800.

Attachment restrictions for Copy Services

This section describes attachment restrictions for the standard FlashCopy, FlashCopy SE, and Remote Mirror and Copy features. The most important difference between the standard FlashCopy and FlashCopy SE features is space efficiency.

When you copy a source volume to a target volume with the FlashCopy or Remote Mirror and Copy feature, 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.

Novell NetWare restriction: 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 or Remote Mirror and Copy operation only under the following conditions:

- For AIX and Sun, when the host system is not using a logical volume manager (LVM).
- For AIX, when the host system is using an LVM with `recreatevg` command support.
- For HP using a Fibre Channel connection, when an LVM accesses a host through 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.

Fibre Channel host attachment

This section provides an overview of Fibre Channel attachment to a DS8000, including architecture, topology, and LUN access modes.

Note: The DS8000 does not support direct attachment of Fibre Channel over Ethernet (FCoE); attachment must be made through a Fibre Channel (FC) switch.

Fibre Channel architecture

Fibre Channel architecture provides various communication protocols on the storage unit. The units that are interconnected are referred to as *nodes*. Each node has one or more *ports*.

A storage unit is a node in a Fibre Channel network. Each port on a storage unit Fibre Channel host adapter is a Fibre Channel port. A host is also a node in a Fibre Channel network. Each port attaches to a serial-transmission medium that provides duplex communication with the node at the other end of the medium.

Storage unit architecture supports three basic interconnection topologies or network structures:

- Arbitrated loop
- Point-to-point
- Switched-fabric

Note: IBM supports the arbitrated loop, point-to-point, and switched-fabric topologies. To change the topology, configure the port.

Arbitrated loop topology

Fibre Channel Arbitrated Loop (FC-AL) is a ring topology that enables you to interconnect a set of nodes. The maximum number of ports that you can have on an FC-AL is 127.

The storage unit supports FC-AL as a private loop. It does not support the fabric-switching functions in FC-AL.

The storage unit supports up to 127 hosts or devices on a loop. However, the loop goes through a loop initialization process (LIP) whenever you add or remove a host or device from the loop. LIP disrupts any I/O operations currently in progress. For this reason, you must have only a single host and a single storage unit on any loop.

Note: The storage unit does not support FC-AL topology on adapters that are configured for FICON protocol.

Figure 2 shows an illustration of an arbitrated loop topology configuration that includes two host systems and two storage units.

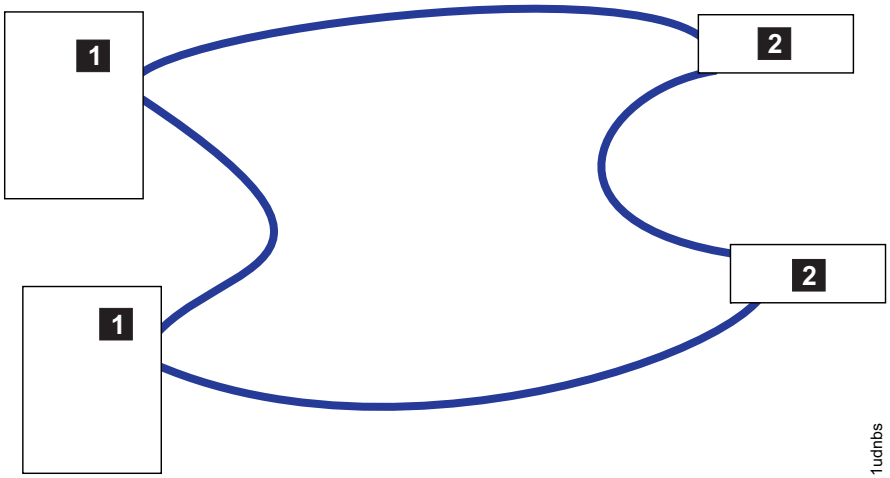

Legend

- 1– is the host system.
- 2– is the storage unit.

Figure 2. Arbitrated loop topology example

Point-to-point topology

The point-to-point topology, also known as direct connect, enables you to interconnect ports directly. Figure 3 shows an illustration of a point-to-point topology configuration that includes one host system and one storage unit.

Legend

- 1 is the host system.
- 2 is the storage unit.

Figure 3. Point-to-point topology example

The storage unit supports direct point-to-point topology at the following maximum distances:

- 1 Gb shortwave adapters have a maximum distance of 500 meters (1640 ft)
- 2 Gb shortwave adapters have a maximum distance of 300 meters (900 ft)
- 4 Gb shortwave adapters have a maximum distance of 300 meters (900 ft)
- 8 Gb shortwave adapters have a maximum distance of 150 meters (492 ft)
- 2 Gb longwave adapters have a maximum distance of 10 kilometers (6.2 mi)
- 4 Gb longwave adapters can have a maximum distance of 4 kilometers (2.5 mi) or 10 kilometers (6.2 mi), depending on the hardware configuration.
- 8 Gb longwave adapters can have a maximum of 4 km (2.5 miles) or 10 km (6.2 miles), depending on the hardware configuration

The maximum distances also vary depending on the cable type. There are three basic types of optical cable fibre. The orange colored cables are shortwave, multimode OM2 type cables. The aqua colored multimode cables are OM3 type and are laser optimized. The yellow colored longwave cables are single fibre. The connection speed in Gigabits per second will determine the distance that is allowed.

Table 7. Connection speed and distance by cable type

Cable type	Speed	Distance
OM2	1 Gbps	500 m (1640 ft)
OM3	1 Gbps	500 m (1640 ft)
OM2	2 Gbps	300 m (900 ft)
OM3	2 Gbps	500 m (1640 ft)
OM2	4 Gbps	150 m (492 ft)
OM3	4 Gbps	270 m (886 ft)
OM2	8 Gbps	50 m (164 ft)
OM3	8 Gbps	150 m (492 ft)

The maximum distance for a longwave cables will also vary depending on the speed and the type of optical transducer. Most small form-factor pluggables (SFPs)

can operate at 10 km but must be selected for that distance and consistent with the connection speed. For example, 4 Gb longwave adapters can have a maximum of 4 km (2.5 miles) or 10 km (6.2 miles) depending on the hardware configuration.

Switched-fabric topology

The switched-fabric topology provides the underlying structure that enables you to interconnect multiple nodes. The distance can be extended by thousands of miles using routers and other storage area network components.

The storage unit supports increased connectivity with the use of Fibre Channel (SCSI-FCP and FICON) directors. Specific details on status, availability, and configuration options that are supported by the storage unit are available at IBM System Storage DS8000 series.

The storage unit supports the switched-fabric topology with point-to-point protocol. You must configure the storage unit Fibre Channel adapter to operate in point-to-point mode when you connect it to a fabric topology.

Figure 4 shows an illustration of a switched-fabric topology configuration that includes two host systems, two storage units, and three switches.

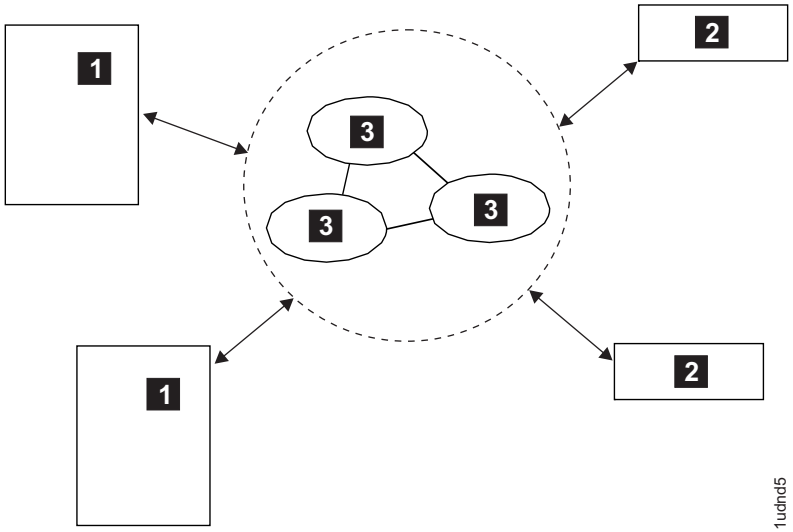

Legend

- 1– is the host system.
- 2– is the storage unit.
- 3– is a switch.

Figure 4. Switched-fabric topology example

Fibre Channel overview for the DS8000 series

Each storage unit Fibre Channel adapter has four or eight ports, and each port has a unique worldwide port name (WWPN). You can configure a port to operate with the SCSI-FCP upper-layer protocol using the DS Storage Manager or the DS CLI. You can add Fibre Channel shortwave and longwave adapters to a DS8000 model.

For details on the host systems that support Fibre Channel adapters, go to the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.

Fibre Channel adapters for SCSI-FCP support the following configurations:

- A maximum of 128 host ports on the DS8700
- A maximum of 16, 4-port 8 Gbps host adapters on the DS8700
- A maximum of 16 host adapters on the DS8700, Model 941 (2-way) and a maximum of 32 host adapters on the DS8700, Model 941 (4-way), which equates to a maximum of 128 Fibre Channel ports
- A maximum of 4 host adapters on the DS8800, Model 951 (2-way) and a maximum of 16 host adapters on the DS8800, Model 95E (4-way), which equates to a maximum of 128 Fibre Channel ports
- A maximum of 506 logins per Fibre Channel port, which includes host ports and PPRC target and initiator ports
- Access to 63700 LUNs per target (one target per host adapter), depending on host type
- Either arbitrated loop, switched-fabric, or point-to-point topologies

Fibre Channel cables and adapter types

This topic provides information about Fibre Channel cables and adapter types.

A storage unit Fibre Channel adapter and FICON host adapter provides four ports with a standard connector. The cables include a standard connector for attachment to the host system with the following items:

- DS8000 adapter
- The four port FC card for DS8000 provides four ports, each using a duplex LC (Lucent) Connector.
- Lucent connector

See the *IBM System Storage DS8000 Introduction and Planning Guide* for detailed information about Fibre Channel cables and adapter types. This document also includes information about cable features and optional cables.

Node-to-node distances for Fibre Channel attachment

DS8000 supports Fibre Channel adapters for extended node-to-node distances.

See the *IBM System Storage DS8000 Introduction and Planning Guide* for a list of longwave and shortwave adapter cables and their distances.

For Fibre Channel attachment, the maximum distance between the following items is 11 km (6 mi):

- Fabric switches
- Host Fibre Channel port
- Link extenders
- Storage unit Fibre Channel port

The maximum distance might be greater than 11 km (6 mi) when a link extender provides target initiator functions or controller emulation functions.

Do not use link extenders with emulation functions on links over which Remote Mirror and Copy operations are performed. This is because of the additional path delay introduced by these units.

LUN considerations for Fibre Channel attachment

You can use a worldwide port name to associate a LUN for Fibre Channel attachment.

For Fibre Channel attachment, LUNs are associated with the Fibre Channel adapter through the worldwide port name (WWPN) for the host adapter.

In a switched fabric configuration, a single Fibre Channel host adapter can have physical access to multiple Fibre Channel ports on the storage unit. In this case, you can configure the storage unit to allow the host to use some or all of the physically accessible Fibre Channel ports on the storage unit.

The maximum number of LUNs that can be assigned to a host is dependent on the host operating system type.

Table 8 lists the maximum number of LUNs that are supported by some of the more common hosts. For more details, see your operating system documentation.

Table 8. LUNs per host operating system

Host operating system type	Number of LUNs per host
AIX	64000
HP	16000
Linux	Thousands
Sun	Thousands[1]
IBM i	64000
Windows	Thousands

[1] Previously for Sun, all storage unit ports shared a single volume group per host, which limited the maximum volumes per host to 256. Currently for Sun, each storage unit port can use a different volume group (256 volumes per port), allowing many thousands of volumes per host. This amount is limited only by the number of ports on the storage unit that are assigned to the Sun host.

Access-restricted mode and profiles for Fibre Channel attachment

This topic describes the access-restricted mode, which is set during initial configuration, and considerations regarding access profiles.

Access-restricted mode

The Fibre Channel architecture allows any Fibre Channel initiator to access any Fibre Channel device, without access restrictions. However, in some environments this flexibility can represent a security exposure. During initial configuration, IBM sets the access mode for the IBM System Storage DS8000 to the access-restricted mode. The access-restricted mode prevents all Fibre Channel attached host systems that do not have access profiles from accessing volumes that you defined in the storage unit.

Access profiles

Any Fibre Channel attached host system that has an access profile can access only those volumes that are defined in the profile. Depending on the capability of the particular host system, an access profile can contain either up to 256 or up to 4096 volumes.

The setup of an access profile is not apparent to you when you use the IBM System Storage DS Storage Manager to configure the hosts and volumes in the storage unit. The following configuration actions can affect the access profile:

- When you define a new Fibre Channel attached host system in the IBM System Storage DS Storage Manager by specifying its worldwide port name (WWPN), the access profile for that host system is automatically created. Initially the profile is empty. That is, it contains no volumes. In this state, the host cannot access any logical volumes that are already defined in the storage unit.
- When you add new logical volumes to the storage unit, the new volumes are assigned to the host. The new volumes are created and are automatically added to the access profile.
- When you assign volumes to Fibre Channel attached hosts, the volumes are added to the access profile.
- When you remove a Fibre Channel attached host system from the IBM System Storage DS Storage Manager, you delete the host and its access profile.

Fibre Channel storage area networks

A Fibre Channel storage area network (SAN) is a specialized, high-speed network that attaches servers and storage devices.

With a SAN, you can perform an any-to-any connection across the network using interconnected elements such as routers, gateways, and switches. Using a SAN can eliminate the connection between a server, storage, and the concept that the server effectively owns and manages the storage devices.

The SAN also eliminates any restriction on the amount of data that a server can access. This restriction is limited by the number of storage devices that can be attached to the individual server. Instead, a SAN introduces the flexibility of networking to enable one server or many heterogeneous servers to share a common storage utility. This might comprise many storage devices, including disk, tape, and optical storage. You can also locate the storage utility far from the servers that use it.

Think of ESCON as the first real SAN. It provides connectivity that is commonly found in SANs. However, it is restricted to ESCON hosts and devices.

Fibre Channel SANs, however, provide the capability to interconnect open systems hosts and System z or S/390 hosts. You can map the protocols for attaching open systems hosts and System z or S/390 hosts to the FC-4 layer of the Fibre Channel architecture.

Fibre Channel worldwide port name identification

The DS8000 uses a worldwide port name (WWPN) to uniquely identify a host port that is connected to a storage unit. The WWPN consists of exactly 16 hexadecimal characters (0 - 9 and A - F).

The DS8000 automatically finds the WWPN for its ports on its host adapters when you attach your host system. You must identify the unique WWPN of each port on your host system that is attached to a DS8000. You use those WWPNs when you configure a host attachment either through the DS Storage Manager or the DS CLI.

You can manually locate a unique worldwide port name for your host ports by performing the steps outlined in the chapter for your host in the *IBM System Storage DS8000 Host Systems Attachment Guide*.

N_Port ID Virtualization (NPIV) is a standardized method for virtualization of a Fibre Channel port. An NPIV-capable Fibre Channel HBA can have multiple N_Port IDs, each with a unique identity and worldwide port name (WWPN). The DS8000 supports configurations that use N_Port ID Virtualization (NPIV) on the host bus adapter or SAN switch.

Open-systems host attachment with Fibre Channel adapters

You can attach a DS8000 series to an open-systems host with Fibre Channel adapters.

Fibre Channel is a 2 Gbps, 4 Gbps or 8 Gbps, full-duplex, serial communications technology to interconnect I/O devices and host systems that are separated by tens of kilometers.

The IBM System Storage DS8000 series supports SAN connections of up to 2 Gbps with 2 Gbps host adapters, up to 4 Gbps with 4 Gbps host adapters, and up to 8 Gbps with 8 Gbps host adapters. The DS8000 series negotiates automatically, determining whether it is best to run at a 1 Gbps, 2 Gbps, 4 Gbps link speed, or 8 Gbps8 Gbps link speed. The IBM System Storage DS8000 series detects and operates at the greatest available link speed that is shared by both sides of the system.

Fibre Channel technology transfers information between the sources and the users of the information. This information can include commands, controls, files, graphics, video, and sound. Fibre Channel connections are established between Fibre Channel ports that reside in I/O devices, host systems, and the network that interconnects them. The network consists of elements like switches, bridges, and repeaters that are used to interconnect the Fibre Channel ports.

Host configuration using the DS8000 interfaces

This section describes how to configure host connections using the following interfaces:

- IBM System Storage DS Storage Manager
- IBM System Storage DS Command-Line Interface (CLI)

Host configuration using the DS Storage Manager

You can use the DS Storage Manager to create new host connections, which includes defining host ports, mapping host ports to volume groups and defining I/O ports.

For detailed information on host connections and step-by-step instructions for creating a new host connection, see the IBM System Storage DS8000 Information Center at DS8000 Information Center , or refer to the DS Storage Manager online help.

Host configuration using the DS CLI

You can use the DS CLI to create new host connections, which includes defining host ports, mapping host ports to volume groups and defining I/O ports; and to display connection information.

For the most current information on supported operating systems for the DS CLI and host configuration instructions, see the IBM System Storage DS8000 Information Center at DS8000 Information Center or the *IBM System Storage DS8000 Command-Line Interface User's Guide*.

Chapter 2. Apple Macintosh host attachment

This chapter provides the requirements that you need to verify before you attach an Apple Macintosh host system to a DS8000 with Fibre Channel adapters.

This chapter contains the following sections:

- "Installing an ATTO adapter on a MAC Pro or MAC Server host system running OS X"
- "Installing and updating an ATTO driver on a MAC Pro or MAC Server host system running the MAC OS X"
- "Installing a SNS adapter on a MAC Pro or MAC Server host system running OS X" on page 16
- "Locating the WWPN for a MAC host" on page 16
- "Creating a file system for a MAC host running OS X" on page 17

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

The host type selection is AppleOSX for the Apple server attachment.

Installing an ATTO adapter on a MAC Pro or MAC Server host system running OS X

Complete this task to install an ATTO Fibre Channel adapter card on a MAC Pro or MAC Server host system running the MAC OS X operating system.

About this task

To install an ATTO adapter on a MAC Pro or MAC Server host system running OS X, perform the following steps:

Procedure

1. Shut down the server.
2. Install the ATTO 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.
3. Connect the cable to the storage unit port or fabric switch port which is able to reach the storage unit.
4. Start the server.

Installing and updating an ATTO driver on a MAC Pro or MAC Server host system running the MAC OS X

Complete this task to install or upgrade an ATTO Fibre Channel adapter driver on a MAC Pro or MAC Server host system running the MAC OS X.

About this task

To install or upgrade an ATTO adapter driver, perform the following steps:

Note: When you are connecting through a Fibre Channel switch, configure the ATTO adapter ports and Fibre Channel switch ports to automatically negotiate speed and topology. This enables the correct speed and topology to be configured automatically.

Procedure

1. Download the ATTO driver update from the ATTO Technology, Inc. website.
   - If you are using the 8Gb Celerity HBA, download the 8Gb Celerity for IBM DS Series from the Fibre Channel HBAs for Storage Partners page.
2. Install the adapter driver. For specific adapter driver installation instructions, see the installation instructions provided by your host adapter vendor.
3. Verify that the ATTO adapter driver module is loaded using the host system **kextstat** command, look for the module name containing *ATTO*.
4. Install the ATTO adapter configuration tool.
5. Set each channel NVRAM parameter Port Down Retry Count to *8* in case there are more than two ports or channels. Set it to *30* if only one channel or port is available.

Installing a SNS adapter on a MAC Pro or MAC Server host system running OS X

Complete this task to install a SNS Fibre Channel adapter card on a MAC Pro or MAC Server host system running the MAC OS X operating system.

About this task

Perform the following steps:

Procedure

1. Shut down the server
2. Install the SNS 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.
3. Connect the cable to the storage unit port or fabric switch port, which is able to reach the storage unit.
4. Start the server.

Locating the WWPN for a MAC host

This topic describes how to locate the worldwide port name (WWPN) for an MAC host that is running the OSX operating system.

To locate the WWPN without restarting the host, use the following command: **sudo fibreconfig -l**.

Creating a file system for a MAC host running OS X

This topic describes how to create and use file systems on the storage unit for a MAC host running the OS X operating system.

You can create a file system using either the GUI, or the CLI.

Using the GUI

Perform the following steps:

1. Go to **Utilities > Disk Utility**.
2. From the left panel, locate the disk that you want use. Then, select **Volume Scheme**.
3. Select the volume format and provide a volume name.
4. Click **Apply**.
5. Click **Partition**.

Using the CLI

Perform the following steps:

1. List disks attaching to host system using the **`$diskutil list`** command.
2. Select the disk that you want to use.
3. Select the volume format you want use and provide a file system name.

Chapter 3. Fujitsu PrimePower host attachment

The host attachment procedure for the Fujitsu PrimePower system is the same as the procedure for the Sun host system. Please refer to Chapter 16, “Sun host attachment,” on page 151 for details.

Chapter 4. HP OpenVMS host attachment

This chapter describes how you attach a Hewlett-Packard (HP) AlphaServer or Integrity server running the OpenVMS operating system to a DS8000 with Fibre Channel adapters.

This chapter contains the following sections:

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

In addition, review the following HP documentation:

- For a discussion about working with Fibre Channel devices and recommendations for system parameters, see *Guidelines for OpenVMS Cluster Configurations*.
- For a detailed description of the AlphaServer console commands that are used in this chapter, see the *WWIDMGR User's Manual*. You can find this document on the Alpha Systems Firmware Update CD.
- For Integrity server documentation, go to the HP Technical Support website.

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. DS8000 host I/O ports no longer must be dedicated to an OpenVMS host. However, non-OpenVMS hosts logged in to 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.

About this task

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

Procedure

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 TCPIP 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*/FCA2*xxx* adapter card in an AlphaServer OpenVMS host

Complete this task to install the KGPSA-*xx*/FCA2*xxx* adapter card in a Hewlett-Packard AlphaServer OpenVMS host.

About this task

To install the adapter card, perform the following steps:

Procedure

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*/FCA2*xxx* 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 starting the host.
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/FCA2*xxx* adapters, replace ¨2¨ in the WWN with ¨1¨.

Example

```
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*/FCA2*xxx* host adapter in an AlphaServer OpenVMS host

Complete this task to set the mode for the KGPSA-*xx*/FCA2*xxx* host adapter in a Hewlett-Packard AlphaServer OpenVMS host.

Before you begin

You must install the KGPSA-*xx*/FCA2*xxx* host adapter before you can set the mode.

About this task

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

Procedure

1. Place the AlphaServer OpenVMS host into console mode.

2. The KGPSA-*xx*/FCA2*xxx* 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/FCA2*xxx* 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.

About this task

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,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 fg*adapter*0**, 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.

About this task

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 26.

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 fg*adapter*0**, 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 topic describes how you can use the OpenVMS Extensible Firmware Interface (EFI) to manage devices for the Hewlett-Packard Integrity server.

About this task

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 5, 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 (fs*x*) 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 5. 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\vms\vms_loader.efi "OpenVMS on $1$DGA78: 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\vms\vms_loader.efi "OpenVMS on $1$DGA78: FGA0.5005-0763-0314-8020"
04. Acpi(HWP0002,100)/Pci(1|0)/Scsi(Pun0,Lun0)/HD(Part1,Sig76D23A51-9B8B-11DB-A618-AA000400FEFF)/\efi\vms\vms_loader.efi "DKA0 PKA0.0" OPT
05. Acpi(HWP0002,100)/Pci(1|0)/Scsi(Pun0,Lun0)/HD(Part1,Sig76D23A51-9B8B-11DB-A618-AA000400FEFF)/\efi\vms\vms_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\vms>
```

Confirming Fibre Channel switch connectivity for OpenVMS

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

Procedure

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:

Example

```
snj2109f16h4:osl> 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:osl>
```

Defining a UDID for OpenVMS storage unit volumes

Each OpenVMS fibre-attached volume requires a user-defined identifier or unit device identifier (UDID). This topic 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 1GGA<UDID>, not as 1DGA<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 DS8000 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 27.

Accessing the storage unit volumes from an OpenVMS host

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

About this task

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

Procedure

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 DS8000 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 DS8000 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 topic 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. For information on using the DS Storage Manager to verify and modify LUN number assignments, see the IBM System Storage DS8000 Information Center at DS8000 Information Center

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 topic apply only if your DS8000 licensed machine code has a version before 5.0.4.0.

Note: You can display the versions of the command-line interface, storage management console (HMC), and licensed machine code by using the DS CLI **`ver -l`** command. This command does not display the version number of the Graphical User Interface (GUI).

The following restrictions are required for the DS8000 host adapter to maintain compatibility with the OpenVMS host system. Compatibility is enabled on a DS8000 adapter port after a defined host establishes Fibre Channel connectivity to one or more storage unit volumes.

- You must dedicate DS8000 adapter ports for only the OpenVMS type. It is recommended that each OpenVMS host adapter be in a fabric zone with one or more DS8000 adapter ports.
- All DS8000 adapter ports that are in a fabric zone with an OpenVMS host adapter must have at least one DS8000 volume assigned to it. This can be the LUN 0 volume.

- Multiple OpenVMS systems can access the same DS8000 adapter port. However, you must define each system for that specific DS8000 port and assign at least one DS8000 volume.
- To re-enable compatibility, you can force the defined OpenVMS host to reestablish connectivity to the DS8000 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.

Chapter 5. HP Tru64 UNIX host attachment

This chapter describes how you attach a Hewlett-Packard (HP) AlphaServer running the Tru64 UNIX operating system to a DS8000 with Fibre Channel adapters.

This chapter contains the following sections:

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

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.

Before you begin

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.

Procedure

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*/FCA2*xxx* adapter card in a Tru64 UNIX host system

Complete this task to install the KGPSA-*xx*/FCA2*xxx* adapter card in a Hewlett-Packard Tru64 UNIX host system.

About this task

To install the adapter card, perform the following steps:

Procedure

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*/FCA2*xxx* 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”.

Example

```
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*/FCA2*xxx* host adapter in a Tru64 UNIX host system

Complete this task to set the mode for the KGPSA-*xx*/FCA2*xxx* host adapter in a Hewlett-Packard Tru64 UNIX host system.

Before you begin

You must install the KGPSA-*xx*/FCA2*xxx* host adapter before you can set the mode.

About this task

The default KGPSA/FCA mode setting is FABRIC, so that directly attaching the AlphaServer to the storage unit using Fibre Channel KGPSA-*xx*/FCA2*xxx* 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:

Procedure

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.

Procedure

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**.

Example

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

This section describes 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.

About this task

To confirm Fibre Channel switch connectivity, perform the following steps:

Procedure

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:

Results

```
snj2109f16h4:osl> 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:osl>
```

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.

About this task

To confirm Fibre Channel storage connectivity, perform the following steps:

Procedure

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.

   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 starting the host.

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>>>
```

For an explanation of the UDID, see "Tru64 UNIX UDID hexadecimal representations" for an explanation of the UDID.

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

Tru64 UNIX UDID hexadecimal representations

This topic 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 an HP AlphaServer Tru64 UNIX host.

About this task

Run 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.

To prepare to boot from the storage unit, perform the following steps:

Procedure

1. Run the `wwidmgr -show wwid` command to display the WWIDs of all assigned storage unit volumes.
2. Determine the 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 38.
3. Run the `wwidmgr -quickset -item` *`item_number`* `-unit` *`unit_number`* command to assign a unit number. The *item_number* is the wwidmgr item number, and *unit_number* is the unit number you choose.

   To find the wwidmgr item number, run the `wwidmgr -show wwid` command. The item number is inside the square brackets of the command output.
4. Run the `init` command to reinitialize the server. When you make changes with the wwidmgr command, they do not take effect until the next system initialization.
5. Run 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, run the console `show device` command to display 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.
6. 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 DGA`*`u`*`.*`, 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.

About this task

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.

Procedure

1. 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 9 on page 40 provides the recommended settings for the DS8000 SCSI device.

Table 9. 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

2. Compile by executing the following command: **ddr_config -c** .
3. Confirm that the values were included by issuing the following command: **ddr_config -s disk** ***IBM 2107*** .
4. 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 */ cdisk_to_def = 60; /60* seconds */**

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 5.*x*host system.

About this task

To verify the Fibre Channel attachment to a Tru64 UNIX 5.*x* host, perform the following steps:

Procedure

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 ibm2107 | 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

Results

```
# ./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.

About this task

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: `# disklevel -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*/FCA2*xxx* 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.

About this task

Perform the following steps to prevent Tru64 AdvFS from losing access to DS8000 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.

Procedure

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
   ```

Chapter 6. HP-UX host attachment

This chapter describes how you attach a Hewlett-Packard (HP) server running the HP-UX operating system to a DS8000 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" on page 48
- "Configuring clustering on an HP-UX host" on page 48

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

For additional SDD information, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* at IBM 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. For data availability and I/O performance, use 4-8 paths per LUN configuration. At least two HBA is recommended.
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 worldwide port name (WWPN) for an HP-UX host.

About this task

To locate the WWPN, perform the following steps:

Procedure

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 begin

Before you set the queue depth, you must attach the host system to a DS8000.

About this task

To set the queue depth for all classes of HP-UX, use the following formula:

```
256 ÷ maximum number of LUNs = queue depth
```

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

To update the queue depth by device level, type `scsictl -m queue_depth=21 /dev/rdsk/$dsksf` where `/dev/rdsk/$dsksf` is the device node. To locate the device node, execute the `ioscan` command.

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 HP Documents HA.

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.

Note: HP-UX 11iv3 is not supported with SDD. The native multipath support driver that comes with HP-UX 11iv3 is supported. See your HP documentation for more information.

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/rdsk/c1t0d0
   pvcreate /dev/rdsk/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)* on the HP website.

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 IBM Storage Software SDD.

Chapter 7. IBM Power Systems host attachment

This chapter describes how you attach an IBM Power Systems host system to a DS8000 with Fibre Channel adapters.

This chapter contains the following sections:

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

IBM Power Systems hosts running IBM i

This section describes how you attach an IBM i host system to a DS8000 with Fibre Channel adapters.

Notes:

1. You cannot use the IBM System Storage Multipath Subsystem Device Driver on IBM i hosts that are running the IBM i operating system. The multipathing capability is included in IBM i operating system.
2. You can assign storage unit LUNs to multiple IBM i Fibre Channel adapters through switches, direct connection, or through a fabric. These multiple adapters must all be within the same IBM i operating system LPAR.
3. For the latest IBM i documentation, go to the IBM Systems Information Centers website at publib.boulder.ibm.com/eserver/.

Locating the WWPN for IBM Power Systems hosts running IBM i

Complete this task to locate the worldwide port name (WWPN) for an IBM i host that is running the IBM i operating system.

Before you begin

You can find the 16-digit WWPN for an IBM i host using either of the following methods:

- Append the number 1000 to the beginning of the 12-digit IEEE address found on the tailstock label or the back side of the Fibre Channel input/output adapter (IOA).
- Display the details for the 2766, 2787, 5749, 5735/5273, 5760, or 5774/5276 IOA Logical Hardware Resource information in the Hardware Service Manager in SST/DST.

About this task

To locate the WWPN, perform the following steps:

Procedure

1. On the IBM i 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 Fibre Channel IOA in the IBM i operating system partition, you must update the WWPN of the new IOA in the storage subsystem, as well as any SAN fabric configurations that you need.

LUN considerations for IBM Power Systems hosts running IBM i

This topic describes several logical unit number (LUN) considerations for IBM i hosts that are running the IBM i operating system.

LUNs are represented in an IBM i environment by a 9-character serial number that is created by the storage unit, for example, 75-1409194. In this example, 75 is the manufacturer's code, 14 is the logical subsystem, 09 is the volume identifier, and 194 is generally the last three digits of the DS8000 worldwide node name (WWNN).

Notes:

1. Fibre Channel attached LUNs are identified as the storage unit device type of 2107 on the IBM i host system.
2. You can specify 1 - 32 LUNs per port on the IBM i Fibre Channel adapter feature 2766, 2787, or 5760.
3. You can specify 1 - 64 LUNs per port on the IBM i Fibre Channel adapter feature 5749, 5774/5276, or 5735/5273.

The following example shows the hardware service manager output of the auxiliary storage hardware resource detail for the 5774 /5276 adapter. Similar information is displayed for the 2766, 2787, 57605749, 5735/5273 and 5760 adapters.

```
Description . . . . . :   Storage IOA
Type-model  . . . . . :   5774-001
Status  . . . . . :   Operational
Serial number . . . . . :   YL10FC63008B
Part number . . . . . :   0000010N7255
Resource name . . . . . :   DC07
Port  . . . . . :   0
Worldwide port name . . . . . :   10000000C9571746
Port  . . . . . :   0
Worldwide port name . . . . . :   10000000C9571745
PCI bus . . . . . :
  System bus  . . . . . :   516
  System board  . . . . . :   0
  System card . . . . . :   0
Storage . . . . . :
  I/O adapter . . . . . :
  I/O bus . . . . . :   127
```

Switch support for IBM Power Systems hosts running IBM i

This topic contains references to supported switches for IBM i hosts that are running the IBM i operating system.

The IBM i operating system supports a homogeneous environment that contains IBM i operating system initiators only. You can establish a homogeneous environment by using the logical zoning of the switch. All hosts within the zone must be IBM i hosts.

For a list of supported switches for IBM i hosts, go to IBM Storage SAN .

Recommended configurations for IBM Power Systems hosts running IBM i

This topic provides information on recommended configurations for IBM i hosts that are running the IBM i operating system.

For the most current information on recommended configurations, go to the IBM Systems Information Centers website at publib.boulder.ibm.com/eserver/.

The following example shows 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                         2107-A82     Operational     DD143
Disk Unit                         2107-A81     Operational     DD140
Disk Unit                         2107-A81     Operational     DD101
```

The following example shows output of the logical hardware resources of the hardware service manager that are associated with the Smart IOA Fibre Channel Adapter:

```
Opt Description                 Type-Model   Status          Resource Name

Virtual IOP                     5774-001     Operational     CMB09
Storage IOA                     5774-001     Operational     DC07
Disk Unit                       2107-A85     Operational     DMP001
```

The following example shows output for the hardware service manager auxiliary storage hardware resource detail for the storage unit:

```
Description.....................:  Disk unit
Type-Model......................:  2107-A85
Status..........................:  Operational
Serial number...................:  50-104F40F
Part number.....................:
Resource name...................:  DMP001
licensed machine code ..........:  FFFFFFFF
Level...........................:  0
PCI bus.........................:
  System bus....................:  516
  System board..................:  0
  System card...................:  0
Storage.........................:
  I/O adapter...................:
  I/O bus.......................:  0
  Controller....................:  2
  Device........................:  1
```

You can define the storage unit LUNs as either protected or unprotected. From a storage unit, physical-configuration viewpoint all IBM i volumes are RAID-5, RAID-6, or RAID-10 volumes and are protected within the storage unit. When you create the IBM i LUNs by using the IBM System Storage DS Storage Manager or DS CLI, you can create them as logically protected or unprotected, depending on the intended use of the LUNs in the IBM i configuration.

Table 10 shows the disk capacity for the protected and unprotected models. Logically unprotecting a storage LUN, allows the IBM i 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 IBM i tool kit, an IBM i 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 10. Capacity and models of disk volumes for IBM i hosts running IBM i operating system

Size	Type	Protected model	Unprotected model
8.5 GB	2107	A01	A81
17.5 GB	2107	A02	A82
35.1 GB	2107	A05	A85
70.5 GB	2107	A04	A84
141.1 GB	2107	A06	A86
282.2 GB	2107	A07	A87

IBM i supports multipath attachment through Fibre Channel as part of the base IBM i operating system support. Obtain a list of supported host adapters on the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC

Note: The FC5708 and FC5729 adapters are not natively supported by IBM i; however they are supported when used in a VIOS partition. DS8000 LUNs can be assigned to the IBM i partition as either vSCSI or vFC (NPIV).

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 DD*xxx* to DMP*xxx* to reflect that it now has more than one path to the disk unit. No user changes are required on the IBM i host to use the new path. Multipath connections can be directly connected or connected through a fabric.

Note: With IBM i operating system 6.1 or later, multipath attachment is supported on an external disk unit that acts as a load source for the IBM i host.

To activate multipathing on the IBM i host, use the IBM System Storage DS Storage Manager.

With IBM i operating system, path information is available for the disk unit hardware configuration from the IBM i Navigator.

To improve the availability of Fibre Channel disks (when there is more than one 2766, 2787, 5749, 5760, 5735/5273 or 5774/5276 I/O adapter), use the following configuration guidelines to help you achieve optimum IBM i I/O performance:

- Use the Disk Magic Analysis tool to calculate the smallest LUN size possible to allow more IBM i-active I/O, without affecting device response times.
- Use RAID-10 LUNs for performance-critical applications.

IBM Power Systems hosts running AIX host attachment

This section describes how you attach an IBM Power Systems running AIX host system to a DS8000 with Fibre Channel adapters.

For additional SDD information, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* at IBM 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 IBM System Storage Multipath Subsystem Device Driver (SDD) supports AIX hosts in a clustering environment. To have failover protection on an open system, SDD requires a minimum of two paths and allows a maximum of 32 paths.
4. For an AIX host, you can use either the switched fabric or arbitrated loop topology. AIX hosts do not support more than one host bus adapter on the loop.

Installing the host attachment package on IBM Power Systems hosts running AIX (from the CD)

Complete this task to install the host attachment package on a System p® AIX host from the CD.

Before you begin

Before installing the host attachment package:

- You must have superuser authority and administrator knowledge to complete the instructions.
- You must have knowledge of the System Management Interface Tool (SMIT).
- Your host system must be attached to a DS8000.

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 a System p AIX host, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* at IBM Storage Software SDD.

About this task

To install the host attachment package on a System p AIX host, complete the following steps:

Procedure

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. Select **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. Select **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 indicates when the process is complete.
13. Press **F10** when the installation process is complete to exit SMIT.
14. Remove the CD.
15. Restart the host.

Installing the host attachment package on IBM Power Systems hosts running AIX (from the SDD website)

Complete this task to install the host attachment package on a System p AIX host from the IBM System Storage Multipath Subsystem Device Driver (SDD) website.

Before you begin

Before installing the host attachment package:

- You must have superuser authority and administrator knowledge to complete the instructions.
- You must have knowledge of the System Management Interface Tool (SMIT).
- Your host system must be attached to a DS8000.

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 a System p AIX host, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* at IBM Storage Software SDD.

About this task

To install the host attachment package on a System p AIX host, complete the following steps:

Procedure

1. Turn on the host system and all attachments.
2. Download the host attachment software at IBM Storage Software SDD.
3. Use FTP to copy the software to the applicable host.
4. If the download package is in tar file format, issue the `tar -xvf` *`package_name`*`.tar` command to untar the package, for example:

   ```
   tar -xvf devices.fcp.disk.ibm.rte.tar
   ```

   For more information, see AIX SDD web-downloaded code installation instructions.
5. From your AIX prompt, type `smit install_update` to go directly to the installation panel.
6. Select **Install Software** and press **Enter**.
7. Type a location for the software, for example, `/tmp`.
8. Press **Enter**. The Install and Update from the Latest Available Software window opens.
9. Select **Software to Install** and press **F4**.
10. Select **Software Packages** and press **F7**.
11. Check the default option settings to ensure that they are what you need.
12. Press **Enter** to install the software. SMIT responds with the following question: `Are you sure?`
13. Press **F10** when the installation process is complete to exit SMIT.
14. Restart the host.

Locating the adapters installed on IBM Power Systems hosts running AIX

Complete this task to locate the adapters that are installed on a System p or RS/6000® host running the AIX operating system.

About this task

To locate the installed adapters, perform the following steps:

Procedure

1. Log in as root.
2. Type `lsdev -Cc adapter`.

   The following example output shows that Fibre Channel adapters fcs0 and fcs1 are installed on the host:

```
hubcap> lsdev -Cc adapter
ppa0    Available 01-R1    CHRP IEEE1284 (ECP) Parallel Port Adapter
sa0     Available 01-S1    Standard I/O Serial Port
sa1     Available 01-S2    Standard I/O Serial Port
sa2     Available 01-S3    Standard I/O Serial Port
siokma0 Available 01-K1    Keyboard/Mouse Adapter
ent1    Available 10-68    10/100 Mbps Ethernet PCI Adapter II (1410ff01)
mg20    Available 10-78    GXT130P Graphics Adapter
ent2    Available 10-80    IBM 10/100 Mbps Ethernet PCI Adapter (23100020)
scsi0   Available 10-88    Wide/Ultra-2 SCSI I/O Controller
scsi1   Available 10-89    Wide/Ultra-2 SCSI I/O Controller
sioka0  Available 01-K1-00 Keyboard Adapter
sioma0  Available 01-K1-01 Mouse Adapter
fcs0    Available 10-70    FC Adapter
fcs1    Available 20-60    FC Adapter
ent0    Available 10-60    IBM 10/100 Mbps Ethernet PCI Adapter (23100020)
```

Locating the WWPN for IBM Power Systems hosts running AIX

Complete this task to locate the worldwide port name (WWPN) for an AIX host.

About this task

To locate the WWPN, perform the following steps:

Procedure

1. Log in as root.
2. Type `lscfg -vl fcs`*x*, where *x* is the adapter number.

   The network address is the Fibre Channel adapter port WWPN value.

   Notes:

   a. The lscfg -vl fcsx ROS level identifies the Fibre Channel adapter firmware level.

   b. To display all adapters with the associated WWPN, issue the following SDD or SDDPCM command:

      - If you are using SDD, type `datapath query wwpn`.
      - If you are using SDDPCM, type `pcmpath query wwpn`.

Attaching IBM Power Systems hosts running AIX

This section provides instructions for attaching an IBM Power Systems host running AIX to a storage unit with Fibre Channel adapters. It is assumed that you have properly connected and zoned the host, FC adapters, and SAN fabric with the storage.

Verifying the storage unit configuration on IBM Power Systems hosts running AIX

Complete this task to verify the configuration of the storage unit for Fibre Channel adapters on an IBM Power Systems host running AIX.

About this task

To verify the storage unit configuration, perform the following steps:

Procedure

Type the following command:

```
lsdev -Cc disk | grep 2107
```

Example

The following is example output for a successful configuration:

```
hdisk3 Available 30-68-01 IBM FC 2107
hdisk4 Available 30-68-01 IBM FC 2107
hdisk5 Available 30-68-01 IBM FC 2107
...
...
```

The following is example output from the **lsdev -Cc disk | grep Other** command, which lists all the disks in the systems that are not configured by the SDD or the SDDPCM multiple path driver.

```
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 IBM Power Systems hosts running AIX:

This section provides tips for troubleshooting a DS8000 that is connected to an IBM Power Systems host that is running AIX.

On a heavily loaded system, you might see lower-than-expected performance and errors indicating that the host adapter is 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. To view the current setting, type `lsattr -El` *adapter_name* `-a max_xfer_size`, where *adapter_name* is the name of the adapter that logs the error.
2. To increase the size of the setting, type `chdev -l` *adapter_name* `-P -a max_xfer_size=0x1000000`.

   Note: To view the range of allowable values for the attribute, type `lsattr -Rl` *adapter_name* `-a max_xfer_size`.
3. Restart the host to use the updated setting.

Making SAN changes on IBM Power Systems hosts running AIX

Review this topic if you are making Storage Area Network (SAN) configuration changes on an IBM Power Systems host. Before you change the switch port that is connected to the storage unit or the switch domain ID, complete this task to remove the Fibre Channel adapter.

About this task

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

To remove the Fibre Channel adapter, perform the following steps:

Procedure

1. Stop all I/O that is running to the affected disks.
2. For each host connected to the storage unit, type `umount` *`filesystem_name`* for all file systems to be unmounted.
3. If you are on an SP host, go to step 4; otherwise, go to step 6.
4. For SP environments, type `suspendvsd` to suspend the virtual shared disks with the volume groups.
5. For SP environments, type `stopvsd` to stop the virtual shared disks that are associated with the volume groups.
6. Type `varyonoff` *`VG_name`* for each 2107/242x volume group to be varied off line.
7. Type `rmdev -dl fcs`*`adapter_number`* `-R` to remove the Fibre Channel adapter. The value of *adapter_number* is the number of the Fibre Channel adapter that you want to remove, for example, `rmdev -dl fcs1 -R`.
8. Make all required SAN changes.
9. Type `cfgmgr` to reconfigure the adapter and rediscover the 2107/242x hdisks.
10. Type `varyonvg` *`VG_name`* for each 2107/242x volume group to be varied online.
11. Type `mount` *`filesystem_name`* to remount all file systems from all hosts that are connected to the storage unit.

Support for Fibre Channel boot on IBM Power Systems hosts running AIX

This section describes the prerequisites, considerations, installation mechanisms, and disk configuration types for a Fibre Channel boot operation on an IBM Power Systems host.

You must have the correct level of the firmware, microcode, and operating system for your adapter and host. For the most current support information, go to the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.

Obtain a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter on the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.

For additional SDD information, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide*. You can find this guide and the latest driver download at IBM Storage Software SDD.

Notes:

1. You must get the microcode from an IBM website.
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. Before you configure a System p host for boot support, ensure that you have SDD or SDDPCM host attachment installed on the system.

Fibre Channel boot considerations for IBM Power Systems hosts running AIX

This topic lists Fibre Channel boot considerations for AIX hosts.

Review the following considerations before you connect a Fibre Channel host:

- You must define the root volume group (rootvg) for each host. On a host with non-MPIO SDD installed, a reserve is placed on all disks in the group when a volume group is varied online. For rootvg, the reserve on this group is not released. You cannot share a rootvg volume among hosts even when you shut down all the hosts except one.

  If MPIO SDDPCM is installed on a host, the root volume group supports multiple paths, and the no_reserve policy is supported. You must not share the rootvg volumes of your host with other hosts. Otherwise, the rootvg data of your host might be accidentally damaged by other hosts.
- With a non-MPIO configured host, you can create only one path to the rootvg volume group during the installation.
- Mirror the rootvg volume group so that base functions and reliability of the operating system for the storage unit are not affected.
- If you use multiboot, scanning for bootable disks takes much longer with many accessible storage elements in the Storage Area Network.

Boot device installation and disk configuration on IBM Power Systems hosts running AIX

This topic lists installation mechanisms and configuration methods for boot installation from Fibre Channel SAN DASD on AIX hosts.

You can use the following mechanisms to direct a PSSP boot installation on an AIX host:

- SAN_DISKID, which is derived from the WWPN and LUN ID

  Note: This is the preferred mechanism.
- PVID
- HDISK

You can use the following methods to configure a disk for an AIX host for a boot install from Fibre Channel SAN DASD:

- Mirrored volume
- Alternate volume
- Individual volume

Attaching to multiple IBM Power Systems hosts running AIX without an 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/6000™) clustering solution.

Considerations for multiple IBM Power Systems hosts running AIX without HACMP

This topic lists considerations for attaching a DS8000 to multiple AIX hosts without an HACMP host.

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.

Review the following considerations if you do not install HACMP and are attaching multiple AIX hosts:

- 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 hosts will most likely affect the availability of the other connected devices.
- Multiple hosts 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 2 - 4 hosts.

Saving data on a storage unit when attaching multiple IBM Power Systems hosts running AIX

Complete this task to preserve the data on a storage unit that was previously connected to an IBM Power Systems host.

About this task

Note: The following procedure removes the volume groups from the host. It does not erase the data on the storage unit.

To save data on the storage unit, perform the following steps:

Procedure

1. Stop all I/O that is running to the affected disks.
2. Type `umount` *`filesystem_name`* to unmount all file systems from all hosts that are connected to the storage unit.
3. Type `fsck` *`filesystem_name`* `-y` on each of the file systems on the storage unit to verify the file system integrity.
4. Type `varyoffvg` *`VG_name`* to vary offline all storage unit volume groups from all hosts that are connected to the storage unit.
5. Type `exportvg` *`VG_name`* to remove all storage unit volume groups from all hosts that are connected to the storage unit.
6. Type `lsdev -C -t 2107* -F name | xargs -n1 rmdev -dl` to delete all the physical volumes (hdisks) on each host that is associated with the storage unit.

What to do next

To restore the data that you save, see "Restoring data when attaching multiple IBM Power Systems hosts running AIX" on page 63.

Restoring data when attaching multiple IBM Power Systems hosts running AIX

Complete this task to restore data on the storage unit when you attach multiple IBM Power Systems hosts to the storage unit.

Before you begin

This procedure assumes that you have preserved the data by performing the steps in "Saving data on a storage unit when attaching multiple IBM Power Systems hosts running AIX" on page 62.

About this task

To restore the data, perform the following steps for each host:

Procedure

1. Ensure that the storage unit physical volumes (hdisks) are available.

   Type `lsdev -Cc disk` to display the hdisks on the host.
2. Type `importvg -y` *`volume_group physical_volume`*, where *volume_group* is the name of the volume group, and *physical_volume* 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 of disks required to keep the volume group on-line ?**
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 /`*`file_systems`* where *file_systems* are the applicable file systems you want to mount.

IBM Power Systems hosts running Linux

This section lists considerations and requirements for installing the IBM Multipath Subsystem Device Driver (SDD) on IBM Power Systems hosts running the Linux operating system.

SDD is supported on Linux through Linux RedHat 4 and SUSE SLES 9. SDD is not supported starting with RedHat 5 and SLES 10; Linux Device Mapper is supported instead. Additional information on SDD is available at www.ibm.com/support/docview.wss?rs=540&context=ST52G7&dc=D430&uid=ssg1S4000107&loc=en_US&cs=utf-8&lang=en.

You can install and configure the SDD using the instructions in the *IBM System Storage Multipath Subsystem Device Driver User's Guide*. You can find this guide and the latest driver download at IBM Storage Software SDD.

Before you install the SDD on IBM Power Systems hosts running Linux, review the following considerations:

- Many Linux distributions enable administrators 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 your hosts are attached to the DS8000 and run the SDD. Some drivers that are supplied by IBM, like the SDD, are dependent an 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 for optimum performance. By allowing automatic updates of the kernel, you risk unexpectedly impacting the performance of your hosts.

- 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 go to IBM Storage Software SDD, and perform the following steps:
  1. Click **Download**.
  2. Click **System Storage Multipath Subsystem Device Driver downloads**.
  3. Click **Subsystem Device Driver for Linux**. This will take you to the SLES 10 Device Mapper Multipath Configuration File for DS8000 / DS6000 / ESS.
- The SDD does not support the following environments for the Linux operating system on an AIX 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
- For hardware considerations, see "Hardware requirements for the SDD on IBM Power Systems hosts running Linux."

Hardware requirements for the SDD on IBM Power Systems hosts running Linux

This topic lists hardware requirements for the Subsystem Device Driver (SDD) on IBM Power Systems hosts running the Linux operating system.

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

- An AIX host that is attached to a DS8000 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 Host system models, OS, adapters and switches for more information about the Fibre Channel adapters that you can use on your Linux host.
  2. Obtain a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter on the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.
- 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 DS8000 port or two switch ports subsequently zoned to DS8000 ports.

Preparing to install the SDD on IBM Power Systems hosts running Linux

This topic describes the tasks you must complete before installing the Subsystem Device Driver (SDD) on IBM Power Systems hosts running the Linux operating system.

Note: Before you can install the SDD, you must configure the storage unit for your host and attach the required Fibre Channel adapters.

Configuring the storage unit

Before you install the SDD, make sure that your DS8000 is configured for multiport access for each LUN. 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 the *IBM System Storage DS8000 Introduction and Planning Guide* for more information about configuring the storage unit for the SDD.

Configuring the Fibre Channel adapters for the storage unit

Before you can install the SDD, you must configure the Fibre Channel adapters and the adapter drivers that are attached to your host. Follow the adapter configuration instructions to configure the adapters. See your Linux on System p documentation for details about installing the device driver on your Fibre Channel adapter.

Installing the SDD on IBM Power Systems hosts running Linux

Complete this task to install the Subsystem Device Driver (SDD) on IBM Power Systems hosts running the Linux operating system.

Before you begin

Before you install the SDD, make sure that your DS8000 is configured for multiport access for each LUN. 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 the *IBM System Storage DS8000 Introduction and Planning Guide* for more information about configuring the storage unit for the SDD.

About this task

Note: See "Upgrading the SDD on IBM Power Systems hosts running Linux" on page 66 for information on upgrading the SDD.

To install the SDD, perform the following steps:

Procedure

1. Log onto your host 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 the following command to view all the files in the package:

   ```
   rpm -qpi IBMsdd-1.6.2.0-1.1.ppc64.sles9.rpm1.4.0.0-6.ppc64.sles9.rpm
   ```
7. Type the following command to install the SDD.

   ```
   rpm -iv IBMsdd-1.6.2.0-1.1.ppc64.sles9.rpm1.4.0.0-6.ppc64.sles9.rpm
   ```

   This displays a message like the following message:

   ```
   Preparing for installation ...
   IBMsdd-1.6.2.0-1.11.4.0.0-6
   ```

Upgrading the SDD on IBM Power Systems hosts running Linux

Complete this task to upgrade the Subsystem Device Driver (SDD) on IBM Power Systems hosts running the Linux operating system.

About this task

You must have a preexisting version of the SDD installed on your host before you can perform these steps. If you are installing the SDD for the first time, see "Installing the SDD on IBM Power Systems hosts running Linux" on page 65.

To install the SDD, perform the following steps:

Procedure

1. Log onto your host 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 the following command to view all the files in the package:

   ```
   rpm -qpi IBMsdd-1.6.2.0-1.1.ppc64.sles9.rpm1.4.0.0-6.ppc64.sles9.rpm
   ```
7. Type the following command to upgrade the SDD:

   ```
   rpm -U IBMsdd-1.6.2.0-1.1.ppc64.sles9.rpm1.4.0.0-6.ppc64.sles9.rpm
   ```

   This displays a message like the following message:

   ```
   Preparing for installation ...
   IBMsdd-1.6.2.0-1.11.4.0.0-6
   ```

Verifying the SDD on IBM Power Systems hosts running Linux

Complete this task to verify the Subsystem Device Driver (SDD) installation on IBM Power Systems hosts running the Linux operating system.

About this task

To verify the SDD installation, perform the following steps:

Procedure

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
sdd-mod.o-*xxxy*	/opt/IBMsdd	The SDD device driver file, where *xxx* represents the kernel level of your host system and *y* represents `smp` or `up`.
vpath.conf	/etc	The SDD configuration file.
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.
3. If the installation is successful, issue the `cd /opt/IBMsdd` command, and then issue the `ls -l` command to list all of the installed SDD components. This creates output like 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 issue the `ls -l` command to view the files in the `/opt/IBMsdd/bin` directory. This creates output like 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 fails, the system displays a message like the following message:

```
package IBMsdd is not installed
```

Configuring and using the SDD

This topic describes where you can find information about configuring and using the Subsystem Device Driver (SDD).

See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* at IBM Storage Software SDD for information on configuring and using the SDD.

Chapter 8. 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 DS8000 with Fibre Channel adapters.

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

For additional information, go to the SAN Volume Controller Support website.

Note: You must create a host entry for every port on every SAN Volume Controller node in your cluster. For example, a 4-node SAN Volume Controller requires 16 host entries.

Chapter 9. IBM NAS Gateway 500 host attachment

This chapter describes how you attach an IBM NAS Gateway 500 host system to a DS8000 with Fibre Channel adapters.

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

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

This section provides information on how 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 topics provide instructions for using these methods.

Obtaining NAS Gateway 500 WWPNs using a web browser

This topic provides instructions for obtaining NAS Gateway 500 WWPNs using a web browser.

About this task

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.

Procedure

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.

About this task

To obtain WWPNs, perform the following steps:

Procedure

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*" > vpd.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
```

Chapter 10. IBM System z or IBM S/390 host attachment

This chapter describes how you attach a System z or S/390 host to a DS8000 with a Fibre Channel, ESCON, or FICON adapter.

Note: ESCON adapters are not supported on DS8700, Model 941 or DS8800, Model 951 .

This chapter contains the following sections:

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

ESCON-attached System z and S/390 hosts overview

This section describes how you can attach DS8100 Models 921 and 931 and DS8300 Models 922, 932, 9A2 and 9B2 to ESCON-attached System z and S/390 hosts.

Note: ESCON adapters are not supported on DS8700, Model 941 and DS8800, Model 951.

With ESCON adapters, the storage unit provides the following configurations:

- A maximum of 32 host ports for DS8100 Models 921 and 931 and a maximum of 64 host ports for DS8300 Models 922, 932, 9A2, and 9B2
- A maximum of 64 logical paths per port
- Access to 16 control-unit images (4096 CKD devices) over a single ESCON port on the storage unit
- Zero to 64 ESCON channels; two per ESCON host adapter
- Two ESCON links with each link that supports up to 64 logical paths

  A DS8100 storage unit supports up to 16 host adapters that provide a maximum of 32 ESCON links per machine. A DS8300 storage unit supports up to 32 host adapters that provide a maximum of 64 ESCON links per machine.

Note: ESCON host channels limit the number of devices per channel to 1024. To fully access 4096 devices on a storage unit, it is necessary to connect a minimum of four ESCON host channels to the storage unit. You can access

the devices through a switch to a single storage unit ESCON port. This method exposes four control-unit images (1024 devices) to each host channel.

The FICON bridge card in ESCON director 9032 Model 5 enables a FICON bridge channel to connect to ESCON host adapters in the storage unit. The FICON bridge architecture supports up to 16 384 devices per channel. This architecture 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 System z and S/390 hosts:

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

For the most current information on supported hosts, operating systems, adapters and switches, go to the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.

ESCON connectivity on System z hosts

This section provides a visual model of ESCON connectivity on System z hosts.

Figure 6 on page 75 shows how a storage unit is attached through ESCON links to different central-electronic complexes and logical partitions (LPARs). Figure 6 on page 75 also shows a configuration that is designed for optimum availability. For optimum availability, you can make the ESCON host adapters available through all bays and spread the connections to any single host across both clusters. For optimum performance, you can have at least eight host adapter ports installed on four ESCON host adapter cards in the storage unit. This setup ensures the best performance from the attached systems.

For information about how to use parallel access volumes (PAVs) for System z hosts, see the IBM System Storage DS8000Information Center at DS8000 Information Center.

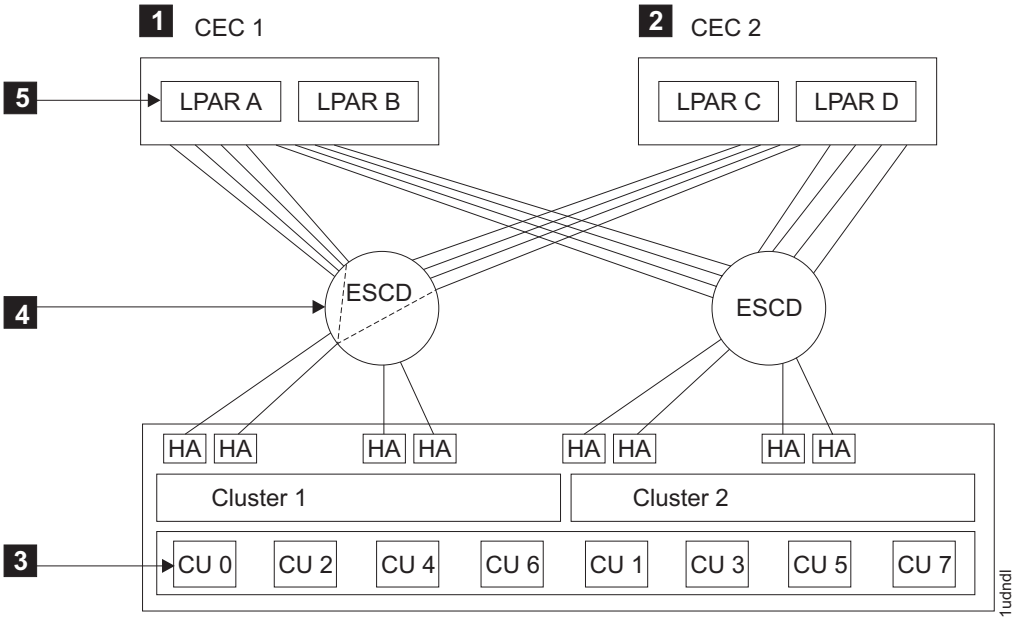

Legend

- 1– computer-electronic complex (CEC) 1
- 2– computer-electronic complex 2
- 3– controller (CU)
- 4– ESCON director (ESCD)
- 5– logically partitioned (LPAR) mode

Figure 6. ESCON connectivity model for System z hosts

Attaching a System z host with ESCON adapters

This section contains information you need to configure a System z host for ESCON attachment to a DS8000.

Controller images and interconnections for ESCON adapters on System z hosts

This topic provides information about controller images on a System z host and lists the types of ESCON interconnections that can be handled by controller images.

For System z host attachments using ESCON, only the first 16 addresses of the logical subsystems (LSSs) can be used.

All controller images are accessible over any installed ESCON physical path. Each controller image can have 1 - 256 devices. The ranges of supported device addresses can be noncontiguous. Devices that are not mapped to a logical device respond and show address exceptions.

Note: When a primary controller connects to a secondary controller, the primary connection converts to a channel.

The controller images can handle the following ESCON interconnections:

- 1 - 256 devices (bases and aliases) per controller image

- 1 - 4096 devices (bases and aliases) with 16 controller images
- 1 - 512 logical paths per controller image
- 1 - 64 logical paths per ESCON port (shared by all controller images)
- 256 path groups per control unit image

Host adapters, cables, distances, and specifications for ESCON adapters on System z hosts

This topic provides information about host adapters, cables, distances, and specifications for ESCON adapters on System z hosts.

Each ESCON host adapter connects to both clusters. A storage unit emulates 0 - 16 of the 3990 logical controllers. Half the logical controllers are in Cluster 1 and half are in Cluster 2. Because the ESCON adapter connects to both clusters, each adapter can address 16 logical controllers. Because each ESCON host adapter provides two host connections, you must order two ESCON cables for each adapter.

The standard distances for ESCON cables are 2 km (1.2 mi) with a 50-micrometer multimode fibre and 3 km (1.9 mi) with 62.5-micrometer multimode fibre. You can extend the distance of the cable to 103 km (64 mi).

Note: For optimum performance, use a cable shorter than 103 km (64 mi).

An ESCON host channel can connect to more than one storage unit port through an ESCON director. The host attaches to one port of an ESCON host adapter in the storage unit. Each storage unit adapter card has two ports.

See the *IBM System Storage DS8000 Introduction and Planning Guide* for a list of the ESCON host adapter features codes. This guide also contains the number of ESCON host adapters, cable group numbers, number of cables, and connector IDs to order for the cables.

Cable lengths and types for ESCON adapters on System z hosts

This topic describes cable lengths and types for ESCON adapters on System z hosts.

All ESCON attachments have a light-emitting diode (LED) interface. The cable that is attached to the host adapter can be up to 2 km (1.24 mi) in length for 50-micrometer multimode fibre, or 3 km (1.86 mi) in length for 62.5-micrometer multimode fibre. You can use the cable to make the connection between the following components:

- Host channel controller
- Peer controller host channel
- Peer controller with appropriate equipment

Note: Appropriate retention hardware to support cable attachments that control bend-radius limits comes with each ESCON host attachment.

Logical paths and path groups for ESCON adapters on System z hosts

This topic provides information about logical paths and path groups for ESCON adapters on System z hosts.

Each ESCON adapter card supports two ESCON ports or links. Each port supports 64 logical paths. With a maximum of 32 ESCON ports, the maximum number of logical paths is 2048.

Each controller image supports up to 64 path groups. Each path group might have up to 8 logical paths. Each controller image supports a maximum of 512 logical paths per control unit image.

Data transfer for ESCON adapters on System z hosts

This topic provides information on data transfer for ESCON adapters on System z hosts.

The ESCON host adapter supports all data input buffer sizes up to 256 bytes. During write operations, the host adapter requests the minimum hexadecimal pacing count of x'02'. For commands whose parameter data length is not determined by the parameter data itself, the full transfer count in the command frame is requested in the first data request. The adapter supports an NDF-R count of 7 (that is, a maximum of eight data requests).

Directors and channel extenders for ESCON adapters on System z hosts

This topic provides information on supported directors and channel extenders for System z hosts.

The storage unit supports IBM ESCON directors 9032 Models 1, 2, 3, and 5. The storage unit supports IBM 9036 channel extenders to the distances that are allowed by the 9036 as described in "Cable lengths and types for ESCON adapters on System z hosts" on page 76. The storage unit supports the 9729 Wavelength Division Multiplexer channel extender up to 50 km (31 mi).

FICON-attached System z hosts overview

This section describes how you can attach the DS8000 storage unit to FICON-attached System z 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 509 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 16 host adapters on DS8700 Model 941 (2-way) and a maximum of 32 host adapters on Model 941 (4-way), which equates to a maximum of 128 host adapter ports
- A maximum of 16, 4-port 8 Gbps host adapters on the DS8700
- A maximum of eight host adapters on the DS8800 Model 951 (4-way), eight ports each, which equates to 64 host adapter ports. With the first expansion model, 95E, another eight host adapters are available, which equates to an additional 64 ports (a maximum of 128 host adapter ports).
- A maximum of 509 logins per Fibre Channel port
- A maximum of 8192 logins per storage unit

- A maximum of 1280 logical paths on each Fibre Channel port
- Access to all 255 control-unit images (8000 CKD devices) over each FICON port
- A maximum of 512 logical paths per control unit image.

Note: FICON host channels limit the number of devices per channel to 16 384. To fully access 65 280 devices on a storage unit, it is necessary to connect a minimum of four FICON host channels to the storage unit. You can access the devices through a switch to a single storage unit FICON port. With this method, you can expose 64 control-unit images (16 384 devices) to each host channel.

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

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

For the most current information on supported hosts, operating systems, adapters and switches, go to the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.

Attaching a System z host with FICON adapters

This topic contains information you need to configure a System z host for FICON attachment to a DS8000.

You can perform FICON channel attachment on the DS8000. See the *IBM System Storage DS8000 Introduction and Planning Guide* for a list of the FICON host adapter feature codes. This guide 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 System z or S/390 hosts

This topic provides an overview of FICON channel attachment on System z or S/390 hosts.

You can use either longwave or shortwave adapters with System z or S/390 hosts. The adapter for the System z has two or four channel ports. The adapter for the S/390 has one channel port.

Note: Depending on your DS8000 model, you might need to purchase a FICON/ESCON license. See the *IBM System Storage DS8000 Introduction and Planning Guide* for more information.

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.

Note: SCSI hosts include System z hosts, which access storage through the SCSI-over-FC protocol (FCP).

When you attach a storage unit to FICON channels through one or more switches, the maximum number of FICON attachments is 509 per storage unit I/O port. The maximum number of logins per storage unit through all I/O ports is 8,192. The directors provide 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 must communicate only with the Fibre Channel devices. The FICON hosts must 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 System z 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 System z 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 System z 800 and 900 model host systems with LIC patches installed and on all other System z models. To support cascaded switches, the fabric must be what is called a high integrity fabric. A high integrity fabric ensures data integrity by detecting any incorrect cabling 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: IBM Storage SAN .

FICON and ESCON considerations for System z and S/390 hosts

This section describes some things to consider before you configure your system with FICON or ESCON adapters for System z hosts.

Setting up ESCON and FICON links

If the system requires *x* ESCON links, where *x* is the number of links to receive the performance and availability attributes that you want, you must consider the number of FICON links that you need. For example, you can map four ESCON links to a single FICON link and maintain equivalent performance.

If the ESCON channel use is low, you can map six or eight ESCON links to a single FICON link. If the ESCON channel use is high, you can map two ESCON links to a single FICON link. ESCON channels can only address the first 16 control unit images on the storage controller, while FICON can address all 255.

Multipathing for ESCON and FICON

Consider the difference between the path groups when you compare FICON to ESCON. For example, for ESCON, you can configure four or eight paths per path group from a host to a storage unit. For ESCON, you can configure at least four paths in the path group to maximize performance.

ESCON is a connection based protocol. When a channel path is busy with a connection for one device, reconnection for other devices must wait for the current connection to complete. With multi-path mode operation, reconnections can occur on any path in a path group. Increasing the number of paths allows reconnections to occur in a more timely basis. If you use eight paths in path groups, you can increase the overall throughput.

FICON allows multiplex of requests for different devices concurrently. There is no problem with reconnections occurring concurrently with other I/O operations on a single path with FICON. The number of paths in the path group depends on the throughput requirement. If it takes *x* paths to satisfy the throughput requirement, where *x* is the number of paths, set the path group to *x*.

Note: *x* must be a minimum of two and cannot exceed a maximum of eight.

Attaching to a FICON channel or a FICON channel-path group

When you attach multiple controllers to a channel, you are connecting serially. You can use a switch (director) for each controller or an ESCON or FICON channel that has a direct connection to the controller. I/O activity does not flow through all the other controllers before you get to the target controller. I/O activity goes directly to the target controller. When multiple controllers are connected to a channel through a switch, you create the logical equivalent of the parallel interconnection.

With the ESCON interface, the channel and controller communicate to form a private connection. None of the other controllers on the channel can communicate with the channel while this private connection is in place. The private connection supports input and output activity between the channel and the controller. It can run slowly, depending upon the factors that affect the controller and the device. The protocol does not allow any of the serially connected controllers to use any spare cycles. The result is poor performance.

FICON does not support a private connection. FICON performs frame (or packet) multiplexing. A configuration with the serially connected controllers communicates with the controllers simultaneously. It can multiplex I/O operations across all controllers simultaneously. No interface cycles are wasted because of a private connection. You can serially connect controllers with FICON without performance degradation.

You must also consider the implications of serially connecting disk drive control units with tape controllers. Tape generally performs much larger I/O operations at any instant in time. Therefore, even with FICON, when you run tape I/O, you can temporarily *lockout* some disk drive I/O operations. It is still better to connect tape and disk drive to different FICON channels.

Identifying the port for TSO commands on System z and S/390 hosts

This topic describes how you can use the **devserv** command to identify ports on a System z or S/390 host.

Figure 7 on page 82 helps you identify the port ID for the MVS™ TSO commands for System z or S/390 hosts. In this figure, the adapter ID numbers B00, B01, B03, and B04 are internal microcode hexadecimal numbers. The strings like cpssfc0010, cpssfc0011, cpssfc0020, and cpssfc0042 are logical adapter interface names for the ports on the adapters.

You need tag numbers for the TSO commands such as `path setup`. To determine the tag number, issue the `devserv` command with the read configuration data (RCD) parameter. The last four characters of a logical adapter interface name form the 2-byte tag number or system adapter (SA) ID.

This side belongs to Cluster 0 (Top Cluster)

This side belongs to Cluster 1 (Bottom Cluster)

	Enclosure 0 (cpssbay00)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B00**	B01	1002	B03	B04	1005
Pt	Logical name					
	cpssfc000	cpssfc001	iss002	cpssfc003	cpssfc004	iss005
0	cpssfc0000	cpssfc0010		cpssfc0030	cpssfc0040	
1	cpssfc0001	cpssfc0011		cpssfc0031	cpssfc0041	
2	cpssfc0002	cpssfc0012		cpssfc0032	cpssfc0042	
3	cpssfc0003	cpssfc0013		cpssfc0033	cpssfc0043	

	Enclosure 1 (cpssbay01)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B06**	B07	1008	B09	B0A	100B
Pt	Logical name					
	cpssfc010	cpssfc011	iss012	cpssfc013	cpssfc014	iss015
0	cpssfc0100	cpssfc0110		cpssfc0130	cpssfc0140	
1	cpssfc0101	cpssfc0111		cpssfc0131	cpssfc0141	
2	cpssfc0102	cpssfc0112		cpssfc0132	cpssfc0142	
3	cpssfc0103	cpssfc0113		cpssfc0133	cpssfc0143	

	Enclosure 2 (cpssbay02)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B0C	B0D	100E	B0F	B10	1011
Pt	Logical name					
	cpssfc020	cpssfc021	iss022	cpssfc023	cpssfc024	iss025
0	cpssfc0200	cpssfc0210		cpssfc0230	cpssfc0240	
1	cpssfc0201	cpssfc0211		cpssfc0231	cpssfc0241	
2	cpssfc0202	cpssfc0212		cpssfc0232	cpssfc0242	
3	cpssfc0203	cpssfc0213		cpssfc0233	cpssfc0243	

	Enclosure 3 (cpssbay03)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B12	B13	B14	B15	B16	1017
Pt	Logical name					
	cpssfc030	cpssfc031	iss032	cpssfc033	cpssfc034	iss035
0	cpssfc0300	cpssfc0310		cpssfc0330	cpssfc0340	
1	cpssfc0301	cpssfc0311		cpssfc0331	cpssfc0341	
2	cpssfc0302	cpssfc0312		cpssfc0332	cpssfc0342	
3	cpssfc0303	cpssfc0313		cpssfc0333	cpssfc0343	

	Enclosure 4 (cpssbay04)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B18	B19	101A	B1B	B1C	101D
Pt	Logical name					
	cpssfc040	cpssfc041	iss042	cpssfc043	cpssfc044	iss045
0	cpssfc0400	cpssfc0410		cpssfc0430	cpssfc0440	
1	cpssfc0401	cpssfc0411		cpssfc0431	cpssfc0441	
2	cpssfc0402	cpssfc0412		cpssfc0432	cpssfc0442	
3	cpssfc0403	cpssfc0413		cpssfc0433	cpssfc0443	

	Enclosure 5 (cpssbay05)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B1E	B1F	1020	B21	B22	1023
Pt	Logical name					
	cpssfc050	cpssfc051	iss052	cpssfc053	cpssfc054	iss055
0	cpssfc0500	cpssfc0510		cpssfc0530	cpssfc0540	
1	cpssfc0501	cpssfc0511		cpssfc0531	cpssfc0541	
2	cpssfc0502	cpssfc0512		cpssfc0532	cpssfc0542	
3	cpssfc0503	cpssfc0513		cpssfc0533	cpssfc0543	

	Enclosure 6 (cpssbay06)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B24	B25	1026	B27	B28	1029
Pt	Logical name					
	cpssfc060	cpssfc061	**iss062**	cpssfc063	cpssfc064	**iss065**
0	cpssfc0600	cpssfc0610		cpssfc0630	cpssfc0640	
1	cpssfc0601	cpssfc0611		cpssfc0631	cpssfc0641	
2	cpssfc0602	cpssfc0612		cpssfc0632	cpssfc0642	
3	cpssfc0603	cpssfc0613		cpssfc0633	cpssfc0643	

	Enclosure 7 (cpssbay07)					
	Slot 0	1	2	3	4	5
	Adapter ID					
	B2A	B2B	102C	B2D	B2E	102F
Pt	Logical name					
	cpssfc070	cpssfc071	iss072	cpssfc073	cpssfc074	iss075
0	cpssfc0700	cpssfc0710		cpssfc0730	cpssfc0740	
1	cpssfc0701	cpssfc0711		cpssfc0731	cpssfc0741	
2	cpssfc0702	cpssfc0712		cpssfc0732	cpssfc0742	
3	cpssfc0703	cpssfc0713		cpssfc0733	cpssfc0743	

For Fibre Channel Adapters:
- Logical names are cpssfcEES (Enclosure, Slot)
- Port logical names are cpssfcEESP (Enclosure, Slot, Port)
- SAID tags are 0xEESP (Enclosure, Slot, Port)

For ESCON Adapters:
- Logical names are “cpssnpEES“ (Enclosure, Slot)
- Port logical names are “cpssnpEESP“ (Enclosure, Slot, Port)
- Adapter IDs are 8XX instead of BXX.

For Device Adapters:
- logical names are “issEES“ (Enclosure, Slot)
- adapter IDs are 10XX.

FC Adapter ID = 0B00 + EE*6 + slot

** **1750** has 2 HAs: B00/cpssfc000 and B06/cpssfc010
and 2 DAs: 1001/iss001 and 1007/iss011

Enclosure numbers (EE) for expansion frames (from the back):

Rack 5		Rack 4		Rack 3		Rack 2		Rack 1	
10 B60	11 B66	C B48	D B4E	8 B30	9 B36	4 B18	5 B1E	0 B00	1 B06
12 B6C	13 B72	E B54	F B5A	A B3C	B B42	6 B24	7 B2A	2 B0C	3 B12

f2c00943

Figure 7. 2107 Adapter Identification Map. A graphical representation of Rack 1 and Rack 2. In the figure, Rack 1 includes enclosures 0, 2, 4, and 6. Rack 2 includes enclosures 1, 3, 5, and 7. Each enclosure includes slot numbers 0, 1, 2, 3, 4, and 5 which run along the top, and port numbers 0, 1, 2, and 3 which run down the left side. Slot numbers 2 and 5 are reserved for host attachment. An adapter ID and logical name are associated with each slot number.

FICON and ESCON migration overview for System z and S/390 hosts

This section describes migration on a System z or S/390 host between FICON and ESCON adapters and between a FICON bridge to a FICON adapter.

Notes:

1. FICON attachment is supported on all 2107/242x models.
2. FICON support consists of hardware enhancements for enterprise servers, host software upgrades, DS8000 LMC, and adapters. The FICON bridge performs the ESCON to FICON conversion.

Migrating from ESCON to native FICON

This section provides information about how to migrate from ESCON to native FICON.

FICON support

FICON is supported on all 2107 models.

Note: FICON support consists of hardware enhancements for enterprise servers, host software upgrades, DS8000 LMC, and adapters. The FICON bridge performs the ESCON to FICON conversion.

Native ESCON configuration

This topic provides a native ESCON configuration example.

Note: ESCON adapters are not supported on DS8700, Model 941 or DS8800, Model 951.

Figure 8 on page 84 shows an example of a native ESCON configuration. The configuration shows a System z or S/390 host with four ESCON channels that are attached to the storage unit through two ESCON directors. The channels are grouped into a channel-path group for multipathing capability to the storage unit ESCON adapters.

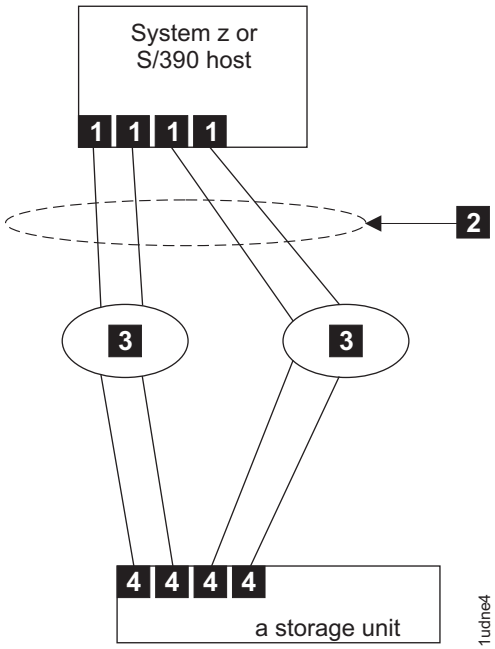

Figure 8. Example of an ESCON configuration on a System z or S/390 host

ESCON to FICON migration example for a System z or S/390 host

This topic provides an ESCON to FICON migration example for a System z or S/390 host.

Figure 9 on page 85 shows another example of a System z or S/390 host with four ESCON channels. In this example, two FICON channels are added to the host. The illustration also shows the channel path group and FICON directors through which the two FICON adapters are installed in the storage unit.

The two FICON directors are not required. You can improve reliability by eliminating a single point of failure. The single point of failure might be present if both FICON channels are connected to a single FICON director. You can connect the FICON channels to the storage unit FICON adapters directly, without directors.

Figure 9 on page 85 also shows four ESCON adapters and two ESCON directors. This configuration gives you the most flexibility for future I/O changes.

Figure 9 on page 85 illustrates the FICON channels that have been added to the existing ESCON channel path group. Because the channel path group has ESCON and FICON channel paths, it makes migrating easier. This intermixing of types of channel-paths enables you to nondisruptively add FICON channels to the host and to add FICON adapters to the storage unit.

Notes:

1. The configuration in Figure 9 on page 85 is supported for migration only.
2. Do not use this configuration for an extended period.
3. Migrate from a mixed channel-path group configuration to an all FICON channel-path group configuration.

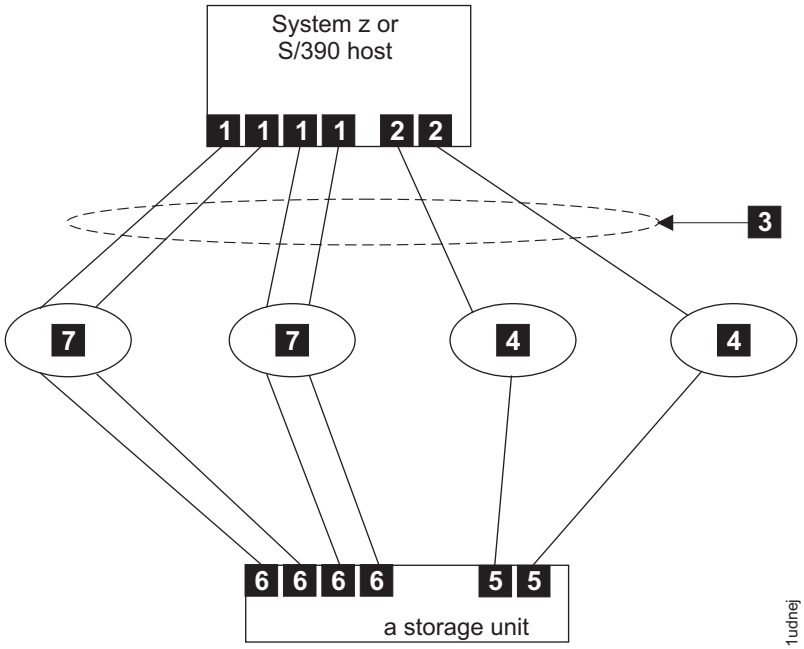

Figure 9. Example of an ESCON to FICON migration on a System z or S/390 host

FICON configuration example for a System z or S/390 host

This topic provides a FICON configuration example for a System z or S/390 host.

Note: ESCON adapters are not supported on DS8700 or DS8800.

Figure 10 on page 86 illustrates how to remove the ESCON paths. The System z or S/390 host has four ESCON channels connected to two ESCON directors. The host also has two FICON channels.

You can remove the ESCON adapters nondisruptively from the storage unit while I/O continues on the FICON paths. You can change the channel-path group definition to include only the FICON director paths to complete the migration to the storage unit with two FICON adapters.

You can retain the ESCON channels on the System z or S/390 host so that you can access other ESCON controllers. You can also keep the ESCON adapters on the storage unit to connect to other System z or S/390 hosts.

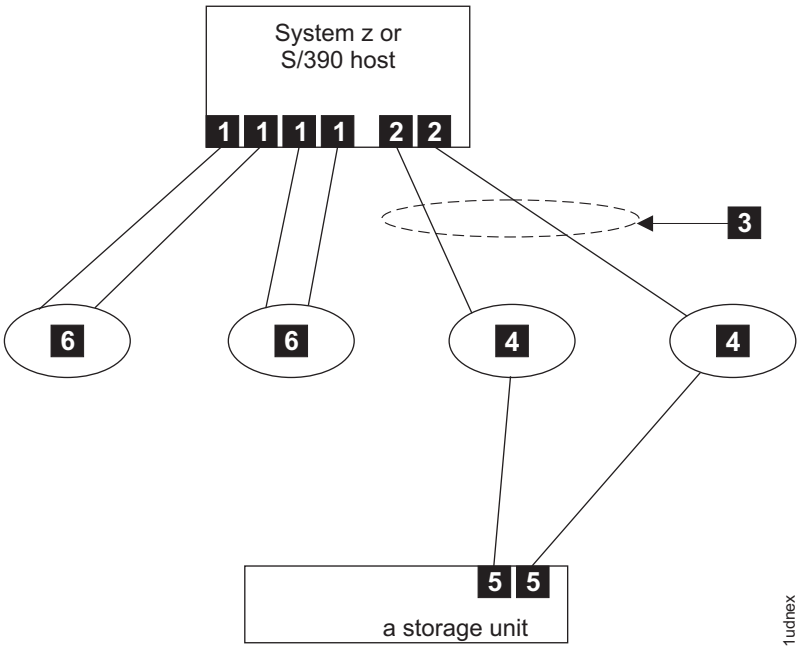

Figure 10. Example of a FICON configuration for a System z or S/390 host

Migrating from a FICON bridge to a native FICON attachment

This section provides examples for migrating from a FICON bridge to a native FICON attachment.

FICON bridge configuration example

This section provides an example of a FICON bridge configuration for a System z or S/390 host.

The FICON bridge is a feature card of the ESCON Director 9032 Model 5. The FICON bridge supports an external FICON attachment and connects internally to a maximum of eight ESCON links. The volume on these ESCON links is multiplexed on the FICON link. You can perform the conversion between ESCON and FICON on the FICON bridge.

Figure 11 on page 87 shows an example of how to configure a FICON bridge. It also shows a System z or S/390 host with two FICON channels attached to two FICON bridges. You can attach the storage unit through the channel-path group to four ESCON links.

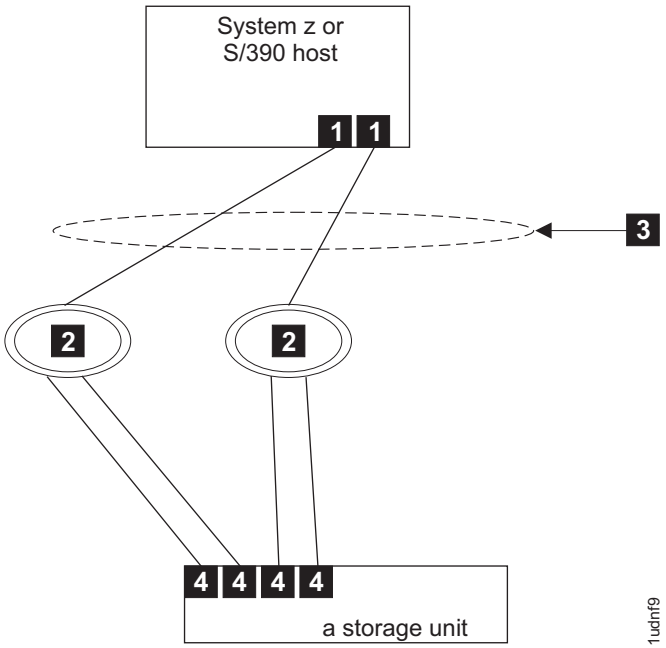

Figure 11. Example of a FICON bridge configuration for a System z or S/390 host

FICON mixed configuration example

This topic provides an example of a FICON mixed configuration for a System z or S/390 host.

Figure 12 on page 88 shows an example of a System z or S/390 host with one FICON channel and one FICON director through a channel-path group and FICON host adapter to the storage unit. Figure 12 on page 88 also shows a System z or S/390 host with one FICON channel and one ESCON director with a FICON bridge through a channel-path group and two ESCON adapters.

Figure 12 on page 88 shows that one FICON bridge was removed from the FICON configuration. The FICON channel that was connected to that bridge is reconnected to the new FICON director. The storage unit FICON adapter is connected to this new director. The channel-path group was changed to include the new FICON path. The channel-path group is a mixed ESCON and FICON path group. I/O operations continue to the storage unit devices across this mixed path group. Access to the storage unit devices is never interrupted because all the actions are nondisruptive.

Notes:

1. The configuration in Figure 12 on page 88 is supported for migration only.
2. Do not use this configuration for an extended time.
3. Migrate from a mixed channel-path group configuration to an all FICON channel-path group configuration.

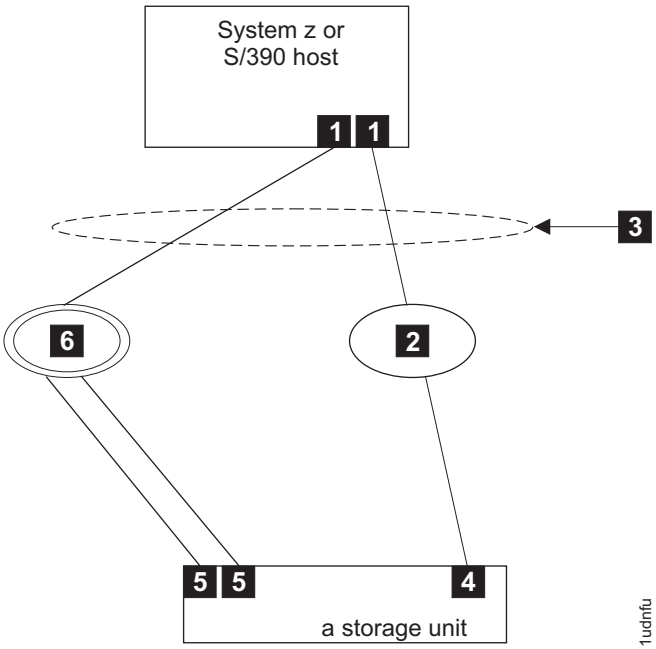

Figure 12. Example of how to add a FICON director and a FICON adapter on a System z or S/390 host

Native FICON configuration on a System z or S/390 host

This topic provides a native FICON configuration example for a System z or S/390 host.

Figure 13 on page 89 shows a System z or S/390 host with two FICON channels connected to two FICON directors through a channel-path group to two FICON adapters. The second bridge has been removed and a FICON director has been added. The channel-path group has only the FICON paths.

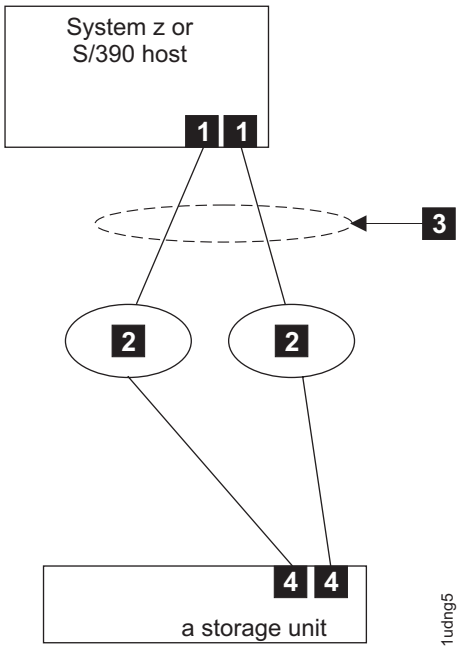

Figure 13. Example of a native FICON configuration on a System z or S/390 host

Linux on System z host attachment

This section describes how to attach the DS8000 storage unit to a System z host running the Linux operating system.

Before you attach a storage unit to any host, review the information in "General requirements for attaching a host" on page 1.

Running Linux on System z hosts

You can run Linux on a System z host in a Logical Partition (LPAR) or as a z/VM guest.

z/VM provides virtualization of central processing units, I/O subsystems, and memory. You can have hundreds of Linux systems running on a single System z host. For example, you can offer a complete Linux server environment to each of the application developers and host production systems that are on the same System z host.

For SUSE Linux Enterprise Server 10 or 11 and Red Hat Enterprise Linux 5.0 or 6.0, IBM supports 64-bit Linux on System z for a storage unit. FICON, ESCON, and Fibre Channel Protocol (FCP) attached hosts are also supported.

Attachment requirements for System z hosts running Linux

This section describes the requirements for attaching the DS8000 to a System z host running the Linux operating system.

Before you begin, complete the following tasks:

- Check the LUN limitations for your host system.

- Ensure that you have the documentation for your host system and the DS8000. You can access DS8000 publications on the Support website at IBM Server Storage Disk or in the DS8000 Information Center at DS8000 Information Center.
- Ensure that you have installed the correct operating systems and are running a supported Linux kernel.
- Review the documentation for your Linux distribution.
- Review the documentation for Linux on System z on the IBM developerWorks® website at IBM Developerworks Linux 390. Check the following information for the Linux kernel that you are using:
  - Supported platforms and environments
  - Hardware and software requirements

Linux distributions for System z hosts

Ensure that you use a Linux distribution that is supported by IBM on your System z host.

IBM supports the following Linux distributions for System z hosts:

- SUSE Linux Enterprise Server
- Red Hat Enterprise Linux AS

HBAs for System z hosts running Linux

Ensure that you use the correct Linux host bus adapter (HBA) and software on your System z host.

Table 11 lists the supported HBAs for System z hosts running the Linux operating system.

For the most current interoperability information about supported HBAs and platform levels, go to Host system models, OS, adapters and switches.

Table 11. Supported HBAs for Linux on System z

Host	Operating systems	Supported HBAs
zSeries® 800	SUSE Linux Enterprise Server and Red Hat Enterprise Linux AS	FICON Express
zSeries 900	SUSE Linux Enterprise Server and Red Hat Enterprise Linux AS	FICON and FICON Express
zSeries 890 and 990	SUSE Linux Enterprise Server and Red Hat Enterprise Linux AS	FICON Express and FICON Express2
System z9® and System z10	SUSE Linux Enterprise Server and Red Hat Enterprise Linux AS	FICON Express, FICON Express2, and FICON Express4
System z196 and System z114	SUSE Linux Enterprise Server and Red Hat Enterprise Linux AS	FICON Express4, FICON Express8, and FICON Express8S
Note: System z HBAs must be ordered as an additional feature for the host system. HBAs are either factory-installed on the host system or installed in an existing system by an IBM service representative.		

Drivers and firmware for System z hosts running Linux

Ensure that you use the correct kernel version and firmware levels for your System z host.

For the most current interoperability information about supported device driver and firmware levels, go to Host system models, OS, adapters and switches.

Installing and configuring HBAs on System z hosts running Linux

This section provides information on installing and configuring host bus adapters (HBAs) on System z hosts running the Linux operating system.

About this task

To check the installation of an HBA and to configure the HBA to work with the DS8000, perform the following steps:

Procedure

1. Ensure that a supported HBA is installed on the System z host. Table 11 on page 90 lists the supported HBAs.
2. Configure the HBA to run in FICON, ESCON, or FCP mode.
3. When using an HBA in FCP mode and multiple guest systems are sharing the same HBA for access to the storage controller, enable the N-Port-ID Virtualization (NPIV) mode of the HBA.

Configuring Linux on System z HBAs in FICON or ESCON mode

You can attach the DS8000 to Linux on System z hosts using FICON or ESCON attachment. Access to FICON or ESCON-attached storage devices is provided by the DASD device driver of the Linux operating system.

Before you begin

For basic information about FICON and ESCON attachment of a System z host, see the following topics:

- "Attaching a System z host with FICON adapters" on page 78
- "Attaching a System z host with ESCON adapters" on page 75

About this task

For more information, see the latest edition of *Linux on System z Device Drivers, Features, and Commands* for the Linux kernel that you are using. You can access this publication on the IBM developerWorks website at IBM Developerworks Linux 390. This website has specific documentation for Red Hat Enterprise Linux and SUSE Linux Enterprise Server distributions, as well as generic information.

Configuring Linux on System z HBAs in FCP mode

You can attach the DS8000 to Linux on System z hosts using FCP attachment.

About this task

Access to FCP-attached storage devices is provided by the sd (SCSI disk) device driver or the zFCP HBA driver of the Linux operating system. For information about FCP connectivity, go to IBM System z connectivity.

Configuring Linux on System z hosts with zFCP attachment

This topic describes the required tasks for configuring FCP connectivity using the zFCP device driver.

Before you begin

You must configure the Linux operating system before you can use your System z host with the DS8000.

Before you configure the host operating system, the following tasks must be completed:

- An IBM sales representative must install the DS8000.
- You must install the appropriate HBAs.

About this task

After the prerequisite tasks are completed, perform the following steps to configure your Linux on System z hosts:

Procedure

1. Zone the host system to the DS8000 on the Fibre Channel SAN.
2. Create the host connection on the DS8000 using the worldwide port names (WWPNs) and map the volumes to the host, as required.

   For information on using the DS CLI to create hosts and map volume, see the *IBM System Storage DS8000 Command-Line Interface User's Guide* or the Command-line interface section of the DS8000 Information Center at DS8000 Information Center. For information on using the DS Storage Manager to create hosts and map volumes, see the Managing section of the DS8000 Information Center at DS8000 Information Center.
3. Configure the Linux operating system for FCP attachment.

   For more information, see the latest edition of *Device Drivers, Features, and Commands* for the Linux kernel that you are using. You can access this publication on the IBM developerWorks website at IBM Developerworks Linux 390. For additional information go to the publication FC-attached SCSI devices with Linux on System z.
4. Create either volumes or disks on your host using the logical volume manager. If you create disks, create file systems on the disks.

   For more information, refer to your host system publications.

Installing multipath support for Linux on System z hosts with zFCP attachment

This topic describes the required tasks for configuring multipath support if you are using the zFCP device driver or want to use Parallel Access Volumes (PAV).

Before you begin

You must install multipath software on all System z hosts that are attached to the DS8000 if you are using the FCP attachment or parallel access volumes (PAV) with RHEL5, RHEL6, SLES10, or SLES11 distributions. SLES10 requires multipathing software for using PAV up to service pack 3. Newer releases of SLES10 service pack 4 might not require multipathing software for using PAV. Check the release notes of the SLES 10 distribution that you are installing.

About this task

On System z hosts running the Linux operating system, the following software provides multipath support:

- Linux kernel 2.6
- Multipath-tools package

With Linux kernel 2.6, the number of paths per physical volume is limited to 8 paths.

For more information about using the logical volume manager or the multipath-tools package, go to IBM Developerworks Linux 390.

SAN boot support on System z hosts running Linux

You can perform an initial program load (IPL) on a SCSI LUN. However, be aware that the boot can fail if there is no multipath support during the boot process.

For more information about using IPL processes with System z hosts, see the latest edition of *How to use FC-attached SCSI devices with Linux on System z*. You can access this publication on the IBM developerWorks website at IBM Developerworks Linux 390.

Defining the number of disks on System z hosts running Linux

When you define the number of disks on System z hosts running the Linux operating system, you are allocating space for configured disks.

About this task

On the Linux operating system, disks are represented as device files. The maximum number of devices depends on your Linux configuration.

DS8000 storage configuration System z hosts running Linux

Each attached DS8000 LUN has a special device file in the Linux directory `/dev`. The maximum number of devices depends on your Linux configuration.

For more information, see the latest edition of *Device Drivers, Features, and Commands* for the Linux kernel that you are using. You can access this publication on the IBM developerWorks website at IBM Developerworks Linux 390.

Known issues and limitations for System z hosts running Linux

There are several restrictions for System z hosts running the Linux operating system.

For the most current information on restrictions, go to the IBM developerWorks website at IBM Developerworks Linux 390.

For additional information, go to the following websites:

- Red Hat website at Red Hat
- Novell at Novell website

Registered state-change notifications on System z or S/390 hosts

This topic provides information about registered state-change notifications (RSCNs) on System z or S/390 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 RSCNs about each other. You can set the zones.

Review the following 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 must be used in this case to support communication from multiple host ports to share device ports.

Analyzing service information messages for System z and S/390 hosts

Service information messages (SIMs) are generated by a storage unit for System z and S/390 hosts. Before installation, you can use the customization work sheets to record 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 unrecoverable error with possible loss of data. This severity level only applies to DASD SIMs.

2 serious
An unrecoverable 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.

Chapter 11. Linux (x86 or x86_64) host attachment

This chapter describes how you attach a Linux host to a DS8000 with Fibre Channel adapters.

This chapter contains the following sections:

Linux Host Attachment Overview

This chapter is meant to provide a general overview of Linux host attachment. Some knowledge of Linux is expected.

IBM supports the two largest commercial Linux distributions, Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES). You may apply the general information in this chapter to other Linux distributions, however IBM may not provide support for them.

IBM supports a variety of HBAs manufactured by QLogic, Emulex and Brocade. This chapter assumes the HBA was previously installed. Use the HBA driver that is included with the distribution.

Prior to Red Hat Enterprise Linux 4 Update 6 and SUSE Linux Enterprise Server 9 SP3, IBM's Subsystem Device Driver was the only multipathing option. For SDD information, go the IBM System Storage website at IBM Storage Software SDD.

For all newer versions of RHEL and SLES, use the native multipathing driver, Device-Mapper-Multipath (DM Multipath). This chapter assumes DM Multipath will be used.

Configuring a QLogic adapter on Linux

This topic describes QLogic adapter settings that can be changed based on the kernel, distribution, and IBM server hardware that you are attaching to the DS8000. 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 Multipath, 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

This topic describes several methods you can use to locate the worldwide port names (WWPN) for a Linux host.

About this task

- **Locating the WWPN without restarting the host**

  To locate the WWPN without restarting the host, you can use any of the following methods:

  – For QLogic or Emulex adapters, you can find the WWPN in the `/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 `qla`*xxxx* for QLogic adapters or `lpfc` for Emulex adapters.

- If you downloaded a QLogic adapter from the QLogic website instead of using the default version that is shipped with RHEL, to find the WWPN run the `cat /proc/scsi/`*adapter_type*`/`*n* command, where *adapter_type* is the host adapter type (for example, `qla2`*xxx* or `qla2300`), and *n* is the host adapter number for your card.
- For a QLogic adapter that is shipped with RHEL 4.*x* or earlier, to find the WWPN run 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. In most cases, the adapter type is `qla2`*xxx*.
- For a QLogic adapter that is shipped with RHEL 5.*x* or later, the adapter WWPN information might not be available in the`/proc/scsi/`*adapter_type*`/`*n* directory because of the migration to sysfs, If so, run the following command:

  `cat /sys/class/scsi_host/host`*n*`/device/fchost\:host`*n*`/port_name`

  where *n* is the host adapter number for your card.

- **Locating the WWPN when restarting the host**

  To locate the WWPN when restarting the host, perform the following steps:

  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 Fibre Channel adapters display. Scroll down to the adapter you want and press `Enter`.
  3. From the **FAST!Util** menu, scroll down and click **Select Host Adapter**.
  4. Scroll up and select **Configuration Settings**, and then 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.

Example

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 or DS CLI.

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 on a Linux host

This section provides information on managing SCSI disks on a Linux host.

Linux SCSI subsystem overview

This topic 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 a Linux host system

This topic describes LUN identification for a 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: 02 Lun: 00
Vendor: IBM Model: 2107900 Rev: .248
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi1 Channel: 00 Id: 02 Lun: 01
Vendor: IBM Model: 2107900 Rev: .248
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi1 Channel: 00 Id: 03 Lun: 00
Vendor: IBM Model: 2107900 Rev: 5.53
Type: Direct-Access ANSI SCSI revision: 05
Host: scsi1 Channel: 00 Id: 03 Lun: 01
Vendor: IBM Model: 2107900 Rev: 5.53
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. See "Rescanning the SAN fabric"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 restarting 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.

Rescanning the SAN fabric

There are several methods that you can use to force the Linux operating system to recognize disks that are added or removed from the fabric. These are the two most common techniques.

1. **Rescan the SAN by restarting the host**

   A bus rescan is automatically performed when restarting the system.
2. **Rescan the SAN by echoing the /sys filesystem**

For Linux 2.6 kernels only, a rescan can be triggered through the /sys interface without having to unload the host adapter driver or restart the system. The following command scans all channels, targets, and LUNs on host H.

```
echo "- - -" > /sys/class/scsi_host/hostH/scan
```

LUN detection procedures

This topic describes how to manually configure LUNs on your 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 that are used for the SCSI mid-layer driver. Figure 14 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 14. Example of a /proc/scsi/scsi file from a Linux host that only configures LUN 0

There are two ways to work around the issue of only having LUN 0 configured:

- Create a script to manually add devices into /proc/scsi/scsi.
- Detect LUNs automatically when the system is started by modifying the initial ram-disk (initrd).

Creating a script to manually add devices into /proc/scsi/scsi

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 15 on page 101 shows part 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 15. Example of a /proc/scsi/scsi file for a Linux host with configured LUNs

Detecting LUNs automatically when the system is started by modifying the initial ram-disk (initrd)

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 when the system is started.

To set the maximum number of disk devices under Linux to properly detect all volumes in Linux 2.4 kernels, 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. Set this value to the respective maximum number of disks the kernel can support, for example, 128 or 256. The same procedure applies to Linux 2.6 kernels except that the parameter is max_luns.

Perform the following steps to set the maximum number of disk devices:

1. Open the /etc/modules.conf file.
2. For Linux 2.4 kernels, add the following line:

   ```
   options scsi_mod max_scsi_luns=n
   ```

   For Linux 2.6 kernels, add the following line:

   ```
   options scsi_mod max_luns=n
   ```

   n is the total number of LUNs to probe.
3. Save the file.
4. To rebuild the ram-disk associated with the current kernel, issue the `mkinitrd` command for your operating system. In the following examples, *kernel* refers to the `uname –r` output that shows the currently running kernel level, for example, `2.4.21-292-smp`.

   For SUSE distributions, issue the following command:

   ```
   cd /boot
   mkinitrd –k vmlinuz-kernel -i initrd-kernel
   ```

   For Red Hat distributions, issue the following command:

```
cd /boot
mkinitrd –v initrd-kernel.img kernel
```

5. Restart the host.
6. Verify that the boot files are correctly configured for the new initrd image in the /boot/grub/menu.lst file.

Linux device naming

This topic 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
```

Note: If the server is running DM Multipath, these SCSI disk devices would not normally be accessed. Use the corresponding multipathed device under /dev/mapper/ instead.

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 problem identification and resolution

This topic 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 restarting 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 restart 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 Linux file system and partitions on the DS8000

To attach a Linux host to a DS8000 you need to create a Linux file system on the DS8000. You can also partition the DS8000 disk for use by your Linux host by using the fdisk utility.

If you choose to partition the storage space on the DS8000, 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 a Linux host

This topic provides the steps you need to assign the system ID to a Linux partition using the **fdisk** command.

Procedure

1. Assign the system partition ID.
2. Write the information to the partition table on the disk.
3. Exit the fdisk program.

Example

Figure 16 on page 106 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 16. Example of assigning a Linux system ID to the partition

Creating a file system for a Linux host

This topic describes how you create and use file systems on the storage unit for a Linux host.

After you partition the disk as described in "Creating a Linux file system and partitions on the DS8000" on page 104, the next step is to create a file system. Figure 17 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 17. Example of creating a file with the mke2fs command

Figure 18 on page 107 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 18. Example of creating a file with the mkfs command

Fibre Channel boot on a Linux host

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 IBM Storage Software SDD.

Performing a Fibre Channel boot with DM Multipath

This topic describes various ways of performing a Fibre Channel boot on a Linux host with DM Multipath.

Note: To perform these tasks, ensure that you have Fibre Channel FC switch zoning, HBA BIOS setting, and experience using the DS8000 command-line interface.

Configuring Red Hat Enterprise Linux (RHEL) 4 Update 6 or later

Perform the following steps:

1. Zone only one path between the server and DS8000 IO port. Use the first host HBA port.
2. Create a host connection using DS CLI or DS Storage Manager.
3. Create a fixed block (FB) volume and map it to the host through a volume group. This FB volume will be used to load the OS.
4. Enable HBA BIOS loading during boot, and select the newly mapped LUN as a boot disk. Refer to the HBA vendor BIOS setup guide for more information.

   Note: Disable the local hard disk through server BIOS.
5. Load RHEL 4 OS to the DS8000 LUN.
6. Build the DM Multipath module to the *initrd* image after the OS boots up with a single path: `<linux:~>#mkinitrd --with=dm-mod --with=dm-multipath --with=dm-round-robin initrdmp.gz '`*`uname -r`*`'`
7. Obtain the boot disk UUID: `<linux:~>#scsi_id -g -u -s /block/sda`
8. Prepare */etc/multipath.conf* as:

```
defaults {
polling_interval 30
failback immediate
no_path_retry 5
rr_min_io 100
path_checker tur
user_friendly_names yes
}
devices {
# DS8000
device {
vendor "IBM"
product "2107900"
path_grouping_policy group_by_serial
}
}
multipaths {
multipath {
wwid 36005076307ffc77d0000000000001050
alias OSROOT
}
}
```

Note: An alias OSROOT is set for boot device to avoid possible device renumbering, which can cause a boot issue.

9. Edit the *initrd* file to add dependency files and configuration file:
   a. Unpack the *initrd* image:
      ```
      <linux host ~>#mkdir temp
      <linux host ~>#cd temp
      <linux host ~>#gzip -dc ../initrdmp.gz | cpio -id
      ```
   b. Copy the files that you need:
      ```
      <linux host ~>#cp /etc/multipath.conf etc/
      <linux host ~>#cp /sbin/dmsetup.static sbin/dmsetup
      <linux host ~>#cp /sbin/multipath.static sbin/multipath
      <linux host ~>#cp /sbin/kpartx.static sbin/kpartx
      <linux host ~>#cp /sbin/scsi_id.static sbin/scsi_id
      <linux host ~>cp -a /etc/udev/rules.d/ etc/udev/
      ```
   c. Edit the *init* file by adding these lines immediately after the drivers finish loading:
      ```
      /sbin/multipath -v2
      /sbin/dmsetup ls --target multipath --exec "/sbin/kpartx -a"
      ```
   d. Repackage the *initrd* image:
      ```
      <linux host ~>#find ./ | cpio -H newc -o > ../initrdmp
      <linux host ~>#gzip initrdmp
      <linux host ~>#cp initrdmp.gz /boot/
      ```
10. Edit *Grub* or *LiLo* to use the new *initrd* image and disk label OSROOT. (For example, `kernel (hd0,1)/boot/vmlinuz root=/dev/disk/by-name/OSROOT2 vga=0x314 splash=silent resume=/dev/disk/by-name/OSROOT1 elevator=cfq showopts`).

    Note: The OSROOT1 represents the first partition of the OSROOT disk.
11. Update the disk label to /etc/fstab:
    ```
    /dev/mapper/OSROOT2 / ext3 default 1 1
    /dev/mapper/OSROOT1 swap swap pri=42 0 0
    ```
12. Enable the DM Multipath service so it loads during system boot: `#chkconfig level 345 multipathd on`.

13. Enable the second path (or all of the paths if you have more than two paths between the host and the DS8000) by zoning, then configure the HBA BIOS to enable boot disk on the second port and reboot the server to verify the setup.

    Note: If you install Linux with LVM2 so you can manage your root file system, you will need to configure LVM configuration file **lvm.conf**to support the DM Multipath device. Copy **lvm.conf** to the *initrd* image. Your **lvm.conf** file should include:

    ```
    filter = [ "a|^/dev/mapper/.*|", "r|.*|" ]
    types = [ "device-mapper", 253 ]
    ```

Configuring Red Hat Enterprise Linux (RHEL) 5 or later

Perform the following steps:

1. Zone all the paths between server and the DS8000.
2. Create a host connection using DS CLI or DS Storage Manager.
3. Create a fixed block (FB) volume and map it to the host through a volume group. This FB volume will be used to load the OS.
4. Enable HBA BIOS loading during boot, and select the newly mapped LUN as a boot disk. Refer to the HBA vendor BIOS setup guide for more information.

   Note: Disable the local hard disk through server BIOS.
5. Load RHEL 5 OS to server CDROM/DVDROM. After you start the RHEL 5 CD/DVD, type `linux mpath` at the prompt, and press **Enter**. This begins the installation with DM Multipath support.
6. When the disk selection page appears, select `mapper/mpath`*x* for the boot disk.
7. After installation, update `/etc/multipath.conf`:

   ```
   defaults {
   polling_interval 30
   failback immediate
   no_path_retry 5
   rr_min_io 100
   path_checker tur
   user_friendly_names yes
   } devices {
   # DS8000 device {
   vendor "IBM"
   product "2107900"
   path_grouping_policy group_by_serial
   }
   }
   ```

Configuring Red Hat Enterprise Linux (RHEL) 6 or later

Perform the following steps:

1. Zone all the paths between server and the DS8000.
2. Create a host connection using DS CLI or DS Storage Manager.
3. Create a fixed block (FB) volume and map it to the host through a volume group. This FB volume will be used to load the OS.
4. Enable HBA BIOS loading during boot, and select the newly mapped LUN as a boot disk. Refer to the HBA vendor BIOS setup guide for more information.

   Note: Disable the local hard disk through server BIOS.

5. Load RHEL 6 OS to server CDROM/DVDROM. After you start the RHEL 6 CD/DVD, select Specialized Storage Devices on the Storage Devices selection page.
6. Select Multipath devices to start the new installation.
7. After installation, update /etc/multipath.conf:

```
defaults {
polling_interval 30
failback immediate
no_path_retry 5
rr_min_io 100
path_checker tur
user_friendly_names yes
} devices {
# DS8000 device {
vendor "IBM"
product "2107900"
path_grouping_policy group_by_serial
}
}
```

Configuring SUSE Linux Enterprise Server (SLES) 9 SP3 or later

Perform the following steps:

1. Zone only one path between the server and the DS8000 IO port. Use the first host HBA port.
2. Create a host connection using DS CLI or DS Storage Manager.
3. Create a fixed block (FB) volume and map it to the host through a volume group. This FB volume will be used to load the OS.
4. Enable HBA BIOS loading during boot, and select the newly mapped LUN as a boot disk. Refer to the HBA vendor BIOS setup guide for more information.

   Note: Disable the local hard disk through server BIOS.
5. Load the SLES 9 OS to this DS8000 LUN.
6. Build DM Multipath module to *initrd* image after OS booting up with single path:
   a. Edit the /etc/sysconfig/kernel file.
   b. Add dm_mod dm_multipath dm_round_robin at the end of line for INITRD_MODULES. (For example, INITRD_MODULES="ata_piix qla2400 dm_mod dm_multipath dm_round_robin reiserfs" .)
   c. Create a new *initrd* image with DM Multipath support:
      <linux:~>#mkinitrd -k /boot/vmlinuz -i initrdmp.gz -x 6000
7. Obtain a boot disk UUID: <linux:~>#scsi_id -g -u -s /block/sda
8. Prepare */etc/multipath.conf* as:

```
defaults {
polling_interval 30
failback immediate
no_path_retry 5
rr_min_io 100
path_checker tur
user_friendly_names yes
} devices {
# DS8000
device {
vendor "IBM"
product "2107900"
path_grouping_policy group_by_serial
}
```

```
}
multipaths {
multipath {
wwid 36005076307ffc77d0000000000001050
alias OSROOT
}
}
```

Note: An alias OSROOT is set for boot device to avoid possible device renumbering, which can cause a boot issue.

9. Adding files need for DM Multipath boot support by performing the following steps:
   a. Unzip the *initrd* image:

```
<linux:~>#mkdir tmp;cd tmp
<linux:~/tmp#cp /boot/initrdmp.gz ./
<linux:~/tmp#gunzip ./initrdmp.gz
```

   b. Mount the *initrd* image to a temporary directory:

```
<linux:~/tmp#mkdir temp
<linux:~/tmp#mount -t ext2 -o loop ./initrdmp temp
<linux:~/tmp#cd temp
```

   c. Copy the DM Multipath related binary, library and configuration files to the *initrd* image:

      Note:

      1) If you do NOT want use LVM to manage your root partition, ignore this step.
      2) If you are using X86_64 SLES9, copy the 64-bit libraries into the lib64 directory.

```
<linux:~/tmp/temp#cp /etc/multipath.conf temp/etc/
<linux:~/tmp/temp#cp /sbin/multipath sbin/
<linux:~/tmp/temp#cp /sbin/dmsetup sbin/
<linux:~/tmp/temp#cp /sbin/kpartx sbin/
<linux:~/tmp/temp#cp /sbin/mpath_* sbin/
<linux:~/tmp/temp#cp /sbin/lvm sbin/
<linux:~/tmp/temp#cp /sbin/vgscan sbin/
<linux:~/tmp/temp#cp /sbin/vgchange sbin/
<linux:~/tmp/temp#cp /sbin/devmap_name sbin/
<linux:~/tmp/temp#cp /bin/awk bin/
<linux:~/tmp/temp#cp /lib/libdevmapper.so.1.01 lib/
<linux:~/tmp/temp#cp lib/libsysfs.so.1 lib/
<linux:~/tmp/temp#mkdir lib/tls
<linux:~/tmp/temp#cp /lib/tls/libc.so.6 lib/tls/
<linux:~/tmp/temp#cp /lib/libselinux.so.1 lib/
<linux:~/tmp/temp#cp /lib/libdl.so.2 lib/
<linux:~/tmp/temp#cp /lib/tls/libm.so.6 lib/tls/
<linux:~/tmp/temp#cp /lib/tls/libc.so.6 lib/tls/
<linux:~/tmp/temp#cp /lib/libsysfs.so.1 lib/
```

   d. Edit the linuxrc file.
   e. Add dm-round-robin driver load line after the HBA driver load line. For example:

```
insmod
/lib/modules/2.6.5-7.308-bigsmp/kernel/drivers/md/dm-round-robin.ko
```

   f. After the last driver-loading line, type:

```
multipath -v2
for i in `dmsetup ls | awk {'print $1'}`;
do
kpartx -a -v /dev/mapper/$i
done
```

g. Repack the *initrd* image:

```
<linux:~/tmp/#umount temp
<linux:~/tmp/#gzip initrdmp
<linux:~/tmp/#mv initrdmp.gz /boot/
```

10. Edit *Grub* or *LiLo* to use the new *initrd* image and disk label OSROOT. (For example, `kernel (hd0,1)/boot/vmlinuz root=/dev/disk/by-name/OSROOT2 vga=0x314 splash=silent resume=/dev/disk/by-name/OSROOT1 elevator=cfq showopts`).

    Note: This step is only needed if root volume is NOT managed by LVM.

11. Change the disk label to /etc/fstab. (For example, `/dev/mapper/OSROOT2 / ext3 default 1`).
12. Instruct the MPIO services to remain persistent upon reboot:

```
# chkconfig multipathd on
# chkconfig boot.multipath on
# chkconfig boot.udev on
```

13. Enable the second path (or the rest of the paths if you have more than two paths between the host and the DS8000) by zoning, then configure the HBA BIOS to enable boot disk on the second port and reboot the server to verify the setup.

Configuring SUSE Linux Enterprise Server (SLES) 10 or later

Perform the following steps:

1. Zone only one path between the server and the DS8000 port. Use the first host HBA port.
2. Create a host connection using DS CLI or DS Storage Manager.
3. Create a fixed block (FB) volume and map it to the host through a volume group. This FB volume will be used to load the OS.
4. Enable HBA BIOS loading during boot, and select the newly mapped LUN as a boot disk. Refer to the HBA vendor BIOS setup guide for more information.

    Note: Disable the local hard disk through server BIOS.

5. Load the SLES 10 OS to the DS8000 LUN.
6. Prepare */etc/multipath.conf* as:

```
defaults {
polling_interval 30
failback immediate
no_path_retry 5
rr_min_io 100
path_checker tur
user_friendly_names yes
}
devices {
# DS8000
device {
vendor "IBM"
product "2107900"
path_grouping_policy group_by_serial
}
}
multipaths {
multipath {
wwid 36005076307ffc77d0000000000001050
alias OSROOT
}
}
```

Note: An alias OSROOT is set for boot device to avoid possible device renumbering, which can cause a boot issue.

7. Update /etc/sysconfig/kernel by adding `module dm_multipath` in INITRD_MODULES. For example:

   ```
   INITRD_MODULES="ata_piix qla2xxx processor thermal fan dm_multipath jbd ext3"
   ```

8. Update /boot/grub/menu.lst to use DM Multipath device during system boot using ONE of the following methods:
   - `kernel /vmlinuz-2.6.16.60-0.21-smp root=/dev/mapper/OSROOT-part3 vga=0x314 resume= /dev/mapper/OSROOT-part2 splash=silent showopts` OR
   - `kernel /vmlinuz-2.6.16.60-0.21-smp root=/dev/mapper/mpatha-part3 vga=0x314 resume= /dev/mapper/mpatha-part2 splash=silent showopts`

   Note: Above name OSROOT-part3 means third partition of the OSROOT disk.

9. Update /etc/fstab to use DM Multipath device using ONE of the following methods:
   - ```
     /dev/mapper/mpatha-part3 / ext3 acl,user_xattr 1 1
     /dev/mapper/mpatha-part1 /boot ext3 acl,user_xattr 1 2
     /dev/mapper/mpatha-part2 swap swap defaults 0 0
     ```
   - OR

     ```
     /dev/mapper/OSROOT-part3 / ext3 acl,user_xattr 1 1
     /dev/mapper/OSROOT-part1 /boot ext3 acl,user_xattr 1 2
     /dev/mapper/OSROOT-part2 swap swap defaults 0 0
     ```

10. Create an *initrd* image to include DM Multipath module support booting from the DM Multipath device: `#mkinitrd -k /boot/vmlinuz -i initrd-2.6.16.60-0.21-smp`.
11. Reboot the server to apply the new *initrd* file.

    Note: If the system cannot boot up, and returns a message such as `can not find /dev/mapper/mpatha-part3`, perform the following steps:

    a. Extract the *initrd* image.

    ```
    #mkdir ~/sanboot; cd ~/sanboot
    #cp /boot/initrd-2.6.16.60-0.21-smp ./initrd.gz
    #gunzip initrd.gz
    #mkdir tmp;cd tmp
    #cpio -id < ../initrd
    ```

    b. Update file *init* to add a 3 second wait after the HBA driver load. For example, if you are using Qlogic:

    ```
    #cd ~/sanboot/tmp
    #vi init
    ```

    Change:

    ```
    echo "Loading qla2xxx"
    modprobe qla2xxx $params
    ```

    To:

    ```
    echo "Loading qla2xxx"
    modprobe qla2xxx $params
    sleep 30
    ```

12. Repack the files into the archive:

    ```
     #cd ~/sanboot/tmp
    #find . | cpio --create --format='newc' > ~/sanboot/newinitrd
    #cd ~/sanboot
    #gzip newinitrd
    #cp newinitrd.gz /boot/initrd-2.6.16.60-0.21-smp
    ```
    ```

Configuring SUSE Linux Enterprise Server (SLES) 11 or later

Perform the following steps:

1. Zone all the paths between server and the DS8000.
2. Create a host connection using DS CLI or DS Storage Manager.
3. Create a fixed block (FB) volume and map it to the host through a volume group. This FB volume will be used to load the OS.
4. Enable HBA BIOS loading during boot, and select the newly mapped LUN as a boot disk. Refer to the HBA vendor BIOS setup guide for more information.

   Note: Disable the local hard disk through server BIOS.
5. Load SLES 11 OS to the DS8000 LUN.
6. On the disk selection page, use *custom partitioning*.
7. On the Expert Partitioner page, select **Hard Disks > Configure > Configure Multipath > Active Multipath**.
8. Partition the DM Multipath managed disk and install SLES 11 on the local hard disk.
9. After installation, update /etc/multipath.conf:

```
defaults {
polling_interval 30
failback immediate
no_path_retry 5
rr_min_io 100
path_checker tur
user_friendly_names yes
}
devices {
# DS8000
device {
vendor "IBM"
product "2107900"
path_grouping_policy group_by_serial
}
}
```

Chapter 12. Microsoft Windows Server host attachment

This chapter describes how you attach a Windows Server 2003 or 2008 host to a DS8000 with Fibre Channel adapters.

This chapter contains the following sections:

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

LUN mapping: When you need access to the data on the target volume, map to the LUNs for the target volumes of the Windows host. This significantly reduces the time it takes for the host system to restart. If you do not map the LUNs, the system can take up to 10 minutes per Remote Mirror and Copy target volume to map the LUNs. If you map the LUNs, perform the LUN mapping after the Remote Mirror and Copy operation and immediately before you access the data. You must restart the host system before you can access the data on the target volume. Perform the mapping when the target volume is mapped to the host through the IBM System Storage DS Storage Manager.

Installing Fibre Channel adapters and adapter drivers on a Windows host

This section describes how you install Fibre Channel adapters and adapter drivers on Windows Server 2003 or 2008.

Installing an Emulex adapter and driver on a Windows host

Complete this task to install and configure an Emulex adapter card and driver on Windows Server 2003 or 2008.

About this task

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

For the most current information on supported Emulex adapters, go to the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.

To install and configure an Emulex adapter and driver, perform the following steps:

Procedure

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 bus adapter driver" on page 2 or www.emulex.com/downloads/ibm/oneconnect-software-kits.html.
3. Connect the cable to the adapter and to the DS8000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS8000, see "Host configuration using the DS Storage Manager" on page 13.

Installing a QLogic adapter and driver on a Windows host

Complete this task to install and configure a QLogic adapter card and driver on Windows Server 2003 or 2008.

About this task

To install and configure a QLogic adapter and driver, perform the following steps:

Procedure

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 bus adapter driver" on page 2 or to support.qlogic.com/support/oem_ibm.asp.
4. Restart the host.

Installing a Brocade adapter and driver on a Windows host

This topic describes where to find information on installing a Brocade adapter and driver on a Windows host.

For information on installing a Brocade adapter and driver on a Windows host, see the Brocade website (http://www.brocade.com/index.page).

Updating a Windows Server device driver

Complete this task to replace an existing Windows Server 2003 or 2008 device driver with an updated version.

About this task

To update a Windows Server device driver, perform the following steps:

Procedure

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 the adapter that you want to update and select **Properties**.
4. Select the **Driver** tab.
5. Click **Update Driver** and follow the instructions that appear on the screen.

Locating the WWPN for a Windows host

Complete this task to locate the worldwide port name (WWPN) for Windows Server 2003 or 2008 with an Emulex or QLogic adapter.

About this task

To locate the WWPN, perform the steps for the adapter you are using.

- **Emulex adapter**
  1. Use the Emulex OneCommand Manager to locate the WWPN. You can download it from the Emulex website at Emulex
  2. After installing OneCommand Manager, to launch the application select **Start** > **Programs** > **Emulex** > **OCManager**.
  3. In the left panel, double-click on the host. It will expand to display the WWPN of the HBA installed on the host.
- **QLogic adapter using the SANsurfer tool**

  Note: This tool enables you to locate the WWPN without restarting the server.
  1. Go to the QLogic website at QLogic and click the **Downloads** tab.
  2. On the Downloads page, select your adapter type, model, and operating system, and then click **Go**.
  3. On the adapter page, go the tools list to download the SANsurfer tool.
- **QLogic adapter using the FAST!Util utility**
  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 Fibre Channel adapters display. Scroll down to the adapter you want and press `Enter`.
  3. From the **FAST!Util** menu:
     a. Scroll down and select **Select Host Adapter**.
     b. Scroll up and select **Configuration Settings**.
     c. Press `Enter`.
  4. From the **Configuration Settings** menu, select **Host Adapter Settings**.
  5. Write down the 16-digit alphanumeric string that is displayed on your screen.

Verifying that the Windows host is configured for storage

Complete this task to determine if Windows Server 2003 or 2008 is configured for storage.

About this task

To verify that the Windows host is configured for storage, perform the following steps:

Procedure

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 that the file was copied to the selected drive.

Installing remote Fibre Channel boot support for a Windows host

Complete this task to install remote Fibre Channel boot support on Windows Server 2003 or 2008.

Before you begin

You must obtain the following items for the Windows host 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 IBM Storage Software SDD.
- The *IBM System Storage DS8000 Introduction and Planning Guide*.
- Either the Windows CD or a set of Windows setup floppy disks with the Windows CD.
- The Windows 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.

About this task

Remote Fibre Channel boot support is available for Windows Server 2003 or 2008 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

Procedure

1. Obtain the WWPN for each host adapter. For instructions, see "Locating the WWPN for a Windows host" on page 117.
2. Zone only one path between the server and DS8000 I/O port.
3. Use the IBM System Storage DS Storage Manager to add the host to the list of host systems that are available for attachment to the storage unit. For detailed information about how to use the IBM System Storage DS Storage Manager, refer to the *IBM System Storage DS8000 Introduction and Planning Guide*.

4. Use the IBM System Storage DS Storage Manager to perform storage allocation to the host system.

Check BIOS level

About this task

Because the host adapter BIOS is responsible for scanning the Fibre Channel bus to locate the remote boot device, ensure 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

Procedure

1. Return to the host system and perform the following steps:
   - **QLogic host adapters**
     a. When the 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 steps:
        – Select **Configuration Settings → Host Bus Adapter Settings**. Set **Host Adapters BIOS** for only one of the adapters to `Enabled` and set **Connection Options** appropriately. For example, set **Connection Options** to 1 for fabric point-to-point. Press `ESC` to return to the previous menu
        – Select **Configuration Settings → Selectable Boot Settings**. Set **Selectable Boot** to `Enabled` and select the appropriate Boot Port Name and LUN that corresponds to the WWPN of the storage unit.
   - **Emulex host adapters:**
     a. When the BIOS banner appears, press `Alt+E` or `Alt+F5` as indicated, depending on the BIOS level. A list of adapters is displayed. Select an adapter by entering the appropriate number, and perform the following steps:
        – Set only one host adapter to `Enabled` and set **Topology** appropriately. For example, set **Topology** to 1 for fabric point-to-point. Press `ESC` to return to the previous menu.
        – Select **Configure Boot Devices** and configure one boot device which includes the boot device, the boot entry WWPN, and the starting LUN.
2. Exit the BIOS utility and save the changes.

Complete the Windows installation

About this task

Apply the Windows Server 2003 or 2008 installation instructions. These instructions include the following steps:

Procedure

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 the message `Press F6 if you want to install a third party SCSI or RAID driver` by pressing `F6`
4. In the Welcome to Setup window, press `Enter`. The message `Set up of Windows now` is displayed.

5. Answer appropriately in the Windows Licensing Agreement window.
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 operating system will reside.

Complete Windows post-installation

Procedure

1. After the Windows Server 2003 or 2008 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` file to install the SDD.
2. Update zoning on the switch to permit the availability of multiple paths.
3. If you are using Windows 2003, install SDD; if you are using Windows 2008, install SDDDSM.

   Note: If running Windows 2008, you must install the MPIO feature on the server before installing SDDDSM.

Chapter 13. VMware ESX Server host attachment

This chapter describes how you attach a VMware ESX Server host to a DS8000 with Fibre Channel adapters.

Note: The information in this chapter refers to VMware ESX Server v2.*x*, v3.*x*, and v4.*x*.

This chapter contains the following sections:

- "Installing the Emulex adapter for a VMware ESX Server host"
- "Installing the QLogic adapter and driver for a VMware ESX Server host" on page 122
- "Defining the number of disk devices on a VMware ESX server" on page 123
- "SCSI disk considerations for a VMware ESX server host" on page 123
- "Configuring a VMware ESX Server host" on page 126
- "Copy Services considerations for VMware ESX servers" on page 129

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

In addition, review the VMware product documentation for the version of VMware that you are using. You can access the VMware documentation at VMware.

LUN limitations: The calculated LUN limitation is a combination of the following VMware limitations:

- Maximum of 64 LUNs running I/Os concurrently
- Less than 1024 LUNs managed by multipathing

Number of paths [1]	**Maximum number of LUNs configured on storage unit / 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 a VMware ESX Server host

This topic describes how you install the Emulex adapter for a VMware ESX Server host.

About this task

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:

Procedure

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 qla*xxxx* 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 www.emulex.com/downloads/ibm/oneconnect-software-kits.html for the latest documentation.

Installing the QLogic adapter and driver for a VMware ESX Server host

Complete this task to install a QLogic adapter and driver for a VMware ESX Server host.

About this task

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 procedure is an example for a QLogic QLA23*xx* adapter. The configuration for your adapter might differ.

Procedure

1. Install the QLogic adapter in the host system.
2. Connect the cable to the adapter and to the DS8000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS8000, see “Host configuration using the DS Storage Manager” on page 13.
3. Restart the server.
4. Install 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 bus adapter driver” on page 2 or support.qlogic.com/support/oem_ibm.asp.
5. Restart the server.

Defining the number of disk devices on a VMware ESX server

This topic 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 a VMware ESX server host

This section describes how you identify LUNs and disk devices, and provides information on persistent binding.

LUN identification for the VMware ESX console

This topic provides LUN identification information for the VMware ESX console.

Each of the attached storage unit LUNs has a special device file in the /dev console operating system directory.

For releases prior to v3.*x*, there is a maximum of 128 Fibre Channel disks entries in the following format:

```
/dev/vsd [0-127]
```

For v3.*x*, the maximum is 256 and the format is `dev/sdnn`, where *n* is either a letter or a space.

From the VMware ESX console, the storage unit can be traced through the following 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 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 (lpfccd for Emulex or qla23*xx* 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 19 on page 124 provides an example of QLogic output and Figure 20 on page 124 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-qla1-adapter-node=200000e08b0c2109;
scsi-qla1-adapter-port=210000e08b0c2109;
scsi-qla1-target-0=5005076300c6a3bb;
scsi-qla1-target-1=200500a0b80f0a15;
scsi-qla1-target-2=5005076300c6a6b8;
scsi-qla1-target-3=5005076300c2a6b8;
scsi-qla1-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 19. 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 20. Example of Emulex output for a VMware ESX server

Vmhba level

The /proc/vmware/scsi directory contains one vmhba*n* directory, where *n* is the host adapter number. The directory contains one entry for each device in the format *scsi_id* : *lun_number*. Figure 21 on page 125 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 21. Example listing of a Vmhba directory for a VMware ESX server

Note: In a multipathing configuration, only one 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 22.

```
[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      20788     936149
        1      -       8837       8581      33869        128         64 0          0
          0          0          0            0      14038     762180

     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      14164     765426          0          0
8728000000
  Total    1000      9005       8668      33915        128         64          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 22. Example of Vmhba entries for a VMware ESX server

Note: The last section displays information on multipathing; refer to the 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 topic 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 a VMware ESX Server host

This section describes how you use the fdisk utility to partition a server disk and how you create a file system a VMware ESX Server hostg.

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 a VMware ESX Server host

This topic provides information about partitioning storage unit disks for a VMware ESX Server host.

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 23 provides an example of different fdisk utility options and Figure 24 on page 128 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 23. 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 24. Example of primary partition on the disk /dev/vsd71 for a VMware ESX server

Creating a file system for a VMware ESX Server host

This topic provides information about creating and using VMFS file system for a VMware ESX Server host.

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 a VMware ESX Server host" on page 126, 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. It is included to ease your task of storage management. It could become a mandatory step if you do not configure persistent binding on the VMware ESX server.

Copy Services considerations for VMware ESX servers

This topic 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 DS8000.
- 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 DS8000 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 DS8000 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 VMware).

- Ensure that you create separate DS8000 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.

Chapter 14. Novell NetWare host attachment

This chapter describes how you attach a Novell NetWare host system to a DS8000 with Fibre Channel adapters.

This chapter contains the following sections:

- "Installing an Emulex adapter and driver for a Novell NetWare host"
- "Installing a QLogic adapter and driver for a Novell NetWare host" on page 132

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

Installing an Emulex adapter and driver for a Novell NetWare host

Complete this task to install an Emulex adapter and driver on a Novell NetWare host.

Before you begin

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.

About this task

To install an Emulex adapter and driver, perform the following steps:

Procedure

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 DS8000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS8000, see "Host configuration using the DS Storage Manager" on page 13.
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 bus adapter driver" on page 2 or www.emulex.com/downloads/ibm/oneconnect-software-kits.html.
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. If you are using SDD or a single path driver, insert the following statement at the beginning of the STARTUP.NCF file: `SET MULTI-PATH SUPPORT=OFF`. If you

are using a Netware Native multipath driver, insert the following statement at the beginning of the STARTUP.NCF file: `SET MULTI-PATH SUPPORT=ON` and `LOAD SCSIHD.CDM AEN`.

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.

About this task

Single- and dual-port Fibre Channel interfaces with the QLogic QLA23*xx* and QLA246*x* adapters support the following public and private loop modes:

- Target
- Public initiator
- Private initiator
- Target and public initiator
- Target and private initiator

To install a QLogic adapter and driver, perform the following steps:

Note: The following procedure is an example of a configuration using the QLogic QLA23*xx* adapter. The configuration for your adapter might differ.

Procedure

1. Install the QLogic adapter in the host system.
2. Connect the cable to the adapter and to the DS8000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS8000, see "Host configuration using the DS Storage Manager" on page 13.
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

Table 12. Recommended settings for the QLogic QLA23xx, and QLA246x adapters for a Novell NetWare host (continued)

Parameters	Recommended settings
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
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	• 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.) **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.

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

Complete this task to install a QLogic adapter driver for a Novell NetWare host.

Procedure

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 bus adapter driver" on page 2 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. If you are using SDD or a single path driver, insert the following statement at the beginning of the STARTUP.NCF file: `SET MULTI-PATH SUPPORT=OFF`. If you are using a Netware Native multipath driver, insert the following statement at the beginning of the STARTUP.NCF file: `SET MULTI-PATH SUPPORT=ON` and `LOAD SCSIHD.CDM AEN`.
    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) or Netware native multipath driver, 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`.

Results

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.

Procedure

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.

Chapter 15. Silicon Graphics host attachment

This chapter describes how you attach a Silicon Graphics (SGI) host system to a DS8000 with Fibre Channel adapters.

This chapter contains the following sections:

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

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 the IRIX operating system manages multipathing.
3. You can also connect a nonclustered SGI host system to the storage unit using the switched fabric topology on optical fibre media, SGI Fibre Channel host bus adapters, and the IBM Fibre Channel 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

Complete this task to check the IRIX operating system version on an SGI host system.

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

Complete this task to install a Fibre Channel adapter card for the SGI host system.

About this task

To install an adapter card, perform the following steps:

Procedure

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.

About this task

To verify the installation of the adapter card, perform the following steps:

Procedure

1. Type `# ls -l /hw/scsi_ctlr` to verify that the command created the links for the controllers.

   In the following example output, 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.

Procedure

1. Type `# cd /var/sysgen/master.d` to go to the required 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 parameter settings

Setting	Topology
0	Arbitrated loop
1	Switched-fabric

Table 13. SGI qlfc_use_connection_mode parameter settings (continued)

Setting	Topology
2	Loop preferred (default)
3	Point-to-point preferred, then loop

Installing an optical cable for SGI in a switched-fabric topology

Complete this task to install an optical cable for SGI in a switched-fabric topology.

About this task

To install an optical cable, perform the following steps:

Procedure

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

Complete this task to install an optical cable for SGI in an arbitrated loop topology.

Procedure

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

Complete this task to confirm switch connectivity for SGI.

Procedure

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:osl> 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 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; ospl3_b1a2
 zone:  SGI_Path2
                SGI02; ospl3_b4a2
 ...
 alias: SGI01   21:00:00:e0:8b:04:d5:39
 alias: SGI02   21:00:00:e0:8b:04:ee:3a
 alias: ospl3_b1a2
                50:05:07:63:00:c3:91:74

 alias: ospl3_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.

Procedure

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

Complete this task to confirm storage connectivity for SGI in a switched-fabric topology

About this task

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.

Procedure

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

Complete this task to confirm storage connectivity for SGI in a Fibre Channel arbitrated loop topology.

About this task

This topic 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.

Procedure

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.

Procedure

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, use the XVM volume manager to configure the failover.conf file so that IRIX points I/O operations to the correct path.

Figure 25 provides an example of the process that you can use to configure the failover.conf file.

```
SGI# /usr/bin/fx -x -d /dev/rdsk/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/rdsk/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 25. 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.

Procedure

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 142 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/scsifo -d** command.

Making a connection through an arbitrated-loop topology

Complete this task to make a connection through an arbitrated-loop topology.

Procedure

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 sc3d0l11 sc4d0l11
GroupB sc3d0l17 sc4d0l17
```

2. To confirm the proper failover configuration, use the IRIX `/sbin/scsifo -d` .

Switching I/O operations between the primary and secondary paths

Complete this task to switch 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

This section describes 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/rdsk` entries instead. The following example shows a naming convention for the disk devices.

```
#/usr/bin/fx -x -d /dev/rdsk/5005076300c003b4/lun1vol/c6p30000
```

Configuring storage in a switched fabric topology

This topic describes the commands you can use to configure storage in a switched fabric topology.

Figure 26 and Figure 27 on page 147 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/rdsk/5005076300c003b4/lun1vol/c6p30000
fx version 6.5, Jul  7, 2000
...opening /dev/rdsk/5005076300c003b4/lun1vol/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/        [exe]rcise/        [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/        [exe]rcise/        [r]epartition/

fx>exit
```

Figure 26. 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/lun1s7/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 27. Example commands for the IRIX switched fabric storage configuration utility

Configuring storage in an arbitrated loop topology

This topic describes the commands you can use to configure storage in an arbitrated loop topology.

Figure 28 on page 148 and Figure 29 on page 149 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/rdsk/dks3d0l17s0

...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/        [exe]rcise/        [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/        [exe]rcise/        [r]epartition/

fx>exit
```

Figure 28. 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/rdsk/dks3d0l17s0"
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 29. Example commands for the IRIX arbitrated loop storage configuration utility

Chapter 16. Sun host attachment

This chapter describes how you attach a Sun host system to a DS8000 with Fibre Channel adapters.

This chapter contains the following sections:

Before you attach a host, review the following information:

- For an overview of host attachment requirements, see "General requirements for attaching a host" on page 1.
- For the most current information on supported hosts, operating systems, adapters, and switches, go to the *System Storage Interoperation Center* (SSIC) website: IBM Storage SSIC.

In addition, go to My Oracle Support! and check the latest recommended patch cluster releases that are associated with your operating system. 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 SDD on a Sun host system running Solaris 8 or 9 in either 32-bit or 64-bit kernel mode, or running Solaris 10 in 64-bit kernel mode.
2. 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 8 or 9.

About this task

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.

To install an AMCC adapter and driver, perform the following steps:

Procedure

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 DS8000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS8000, see "Host configuration using the DS Storage Manager" on page 13.
4. Restart the system.
5. Use the instructions in "Downloading and installing a host bus adapter driver" on page 2 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 8, 9, or 10.

About this task

Note: For all Emulex HBAs (LP7000, LP8000, LP8000S, LP9000, LP9002L, LP9002S, LP9002C, LP9402, LP9802, LP9802DC, LP10000, LP10000DC, LP11000, LP11002, LPe11000, LPe11002, LPe11004, LPe12000, LPe12002, and LPe12004), the following conditions apply:

- The native Emulex "lpfc" driver is only available for Solaris 8 and 9
- For Solaris 8 and 9, the Sun SFS "emlxs" driver is available starting from HBA model LP900x to later models (LP1000x, LP/e1100x, LP/e1200x)
- For Solaris 10, the only driver supported is Sun SFS "emlxs" driver

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:

Note: Starting with Emulex driver for Solaris 6.20f, and Qlogic's Solaris driver 4.17, the sd.conf and st.conf files are no longer used to limit the storage presented to the OS. If you are using that driver, you do not need to modify the sd.conf and st.conf files described in this topic.

Procedure

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 DS8000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS8000, see "Host configuration using the DS Storage Manager" on page 13.
5. Restart the system.

6. Use the instructions in "Downloading and installing a host bus adapter driver" on page 2 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 8, 9, or 10.

About this task

To install a QLogic adapter and driver, perform the following steps:

Note: Starting with Emulex driver for Solaris 6.20f, and Qlogic's Solaris driver 4.17, the sd.conf and st.conf files are no longer used to limit the storage presented to the OS. If you are using that driver, you do not need to modify the sd.conf and st.conf files described in this topic.

Procedure

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 DS8000 or switch port that you configured for the host adapter. If you have not configured host adapter ports on your DS8000, see "Host configuration using the DS Storage Manager" on page 13.
4. Restart the system.
5. Use the instructions in "Downloading and installing a host bus adapter driver" on page 2 to download the QLogic driver for your adapter.
6. Install the QLogic host adapter driver package:
   a. Locate the **qla*xxxx*.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*/qla*xxxx*.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 qla*xxxx*.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 gla*xxxx*

      Note: The **Z** suffix is dropped after the file is uncompressed.

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 8, 9, or 10.

About this task

Attention: 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.

To install a Sun adapter and driver, perform the following steps:

Procedure

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 DS8000 or to the switch port that you have configured for the host adapter. If you have not configured host adapter ports on your DS8000, see "Host configuration using the DS Storage Manager" on page 13.
4. Restart the system.
5. If you are using Solaris 8 or 9, follow the instructions in "Downloading and installing a host bus adapter driver" on page 2 to download the latest Sun driver for your adapter.

   If you are using Solaris 10, you can skip this step because the SAN software is included in the operating system.
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 website.
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 DS8000 adapter port that you connected to in step 3.
14. Repeat step 13 as necessary to ensure that all LUNs are recognized.
15. Type `devfsadm` to scan all attached devices.
16. Type `format` to display all LUNs that were discovered by the Solaris operating system.

Configuring host device drivers on a Sun host

This section describes how to update the configuration files for your adapter and the Solaris device driver.

About this task

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.

For configuration instructions, see the following topics:

- "Configuring AMCC, Emulex, and QLogic host device drivers on a Sun host"
- "Configuring Sun-Emulex and Sun-QLogic host device drivers on a Sun host" on page 164

Configuring AMCC, Emulex, and QLogic host device drivers on a Sun host

Complete this task to update the configuration files for your AMCC, Emulex, or QLogic adapter and the Solaris device driver.

About this task

To configure the host device drivers, perform the following steps:

Note: Starting with Emulex driver for Solaris 6.20f, and Qlogic's Solaris driver 4.17, the sd.conf and st.conf files are no longer used to limit the storage presented to the OS. If you are using that driver, you do not need to modify the sd.conf and st.conf files described in this topic.

Procedure

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/jnic`*xxx*`.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 157.

   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 160.

   For Emulex adapters that are using Sun SFS "emlxs" driver, 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 164.

QLogic adapters

Parameters are in the `/kernel/drv/qla2300.conf` file. For parameter descriptions and recommendations, see "Parameter settings for QLogic QLA23*xx* adapters with driver level v5.05 build 11 on a Sun host" on page 161 or "Parameter settings for QLogic QLA23*xx*, QLA246*x*, and QLE256*x* adapters with a San Surf configuration on a Sun host" on page 162.

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.

   Attention:

   - 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.
   - Do not add duplicate target and LUN pairs.

   Figure 30 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 30. Example of sd.conf file entries for a Sun host

You might want to add lines to probe for additional LUNs or targets, as shown in the following example:

```
name="sd" class="scsi"
        target=0 lun=50;
name="sd" class="scsi"
        target=0 lun=51;
```

Delete any lines that represent lpfc targets or LUNs that are not used.

Figure 31 on page 157 shows the start lpfc auto-generated configuration.

Attention: You will delete anything that you put within this auto-generated section if you issue the `pkgrm` command to remove the lpfc driver package.

```
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 31. Example of a start lpfc auto-generated configuration for a Sun host

5. From the Sun UNIX prompt, type the following commands:
   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.

   Note: Starting with Emulex driver for Solaris 6.20f, and with Qlogic Solaris driver 4.17, the sd.conf and st.conf files are no longer used to limit the storage presented to the operating system. If you are using that driver, you do not need to modify the sd.conf and st.conf files described in this topic.

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 on page 158 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.
FcLinkSpeed	**3** (Default) Specifies the desired Fibre Channel link speed: 0: Default to SEEPROM setting 1: Force 1 gigabit per second 2: Force 2 gigabit per second 3: Auto negotiate link speed
FcLinkUpRecoveryTime	**1000** (Default) 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.

Table 14. Recommended configuration settings for AMCC adapters on a Sun host (continued)

Parameters	Recommended settings
FlogiRetryCount	**3** (Default) Specifies the total number of fabric login (FLOGI) attempts before giving up logging in to a switch. Failure prevents participation on a fabric topology.
JniCreationDelay	**5** (Default) 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.
LunRecoveryInterval	**50** (Default) 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 so, increase the value of this parameter.
PlogiControlSeconds	**30** (Default) 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 so, increase the value of this parameter.
PlogiRetryCount	**5** (Default) Specifies the total number of port login (PLOGI) attempts before failing to log in to an SCSI target.
QfullRetryCount	**5** (Default) 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.
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)

Table 14. Recommended configuration settings for AMCC adapters on a Sun host (continued)

Parameters	Recommended settings
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, LP9802, LP10000, LP10000, LP11000, LP11002, LPe11000, LPe11002, LPe12000, and LPe12002 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.)
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.

Table 15. Recommended configuration settings for Emulex adapters on a Sun host (continued)

Parameters	Recommended settings
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 QLogic QLA23*xx* adapters with driver level v5.05 build 11 on a Sun host

You can use these recommended configuration settings for QLogic QLA2310F, QLA2340, and QLA2342 adapters with driver level v5.05 build 11 on a Sun host.

Table 16 on page 162 describes the configuration parameters and recommended settings for QLogic QLA23*xx* adapters with driver level v5.05 build 11. These settings are for a host that is attached directly to DS8000 or through a fabric switch.

Table 16. Recommended configuration settings for QLogic QLA23xx adapters with driver level v5.05 build 11 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
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-binded devices to the operating system

Parameter settings for QLogic QLA23*xx*, QLA246*x*, and QLE256*x* adapters with a San Surf configuration on a Sun host

You can use these recommended settings for QLogic QLA23*xx*, QLA246*x*, QLE246*x*, QLE256*x* adapters on a Sun host. Adapter QLA23*xx* is for storage units with a San Surf configuration (4.06+ driver). Adapters QLA246*x*, QLE246*x*, and QLE256*x* are for storage units with a San Surf configuration (5.0+ driver).

Table 17 on page 163 describes the configuration parameters and recommended settings. These settings are for a host that is attached directly to DS8000 or through a fabric switch.

Table 17. 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	• 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
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: 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 qla*xxxx*.conf file. Figure 32 on page 164 shows an example of a qla*xxxx*.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 32. Example of a qlaxxxx.conf file for a Sun host

Configuring Sun-Emulex and Sun-QLogic host device drivers on a Sun host

Complete this task to update the configuration files for your Sun-Emulex or Sun-QLogic adapter and the Solaris device driver.

About this task

To configure the host device drivers, perform the following steps:

Procedure

1. Type `cd /kernel/drv` to change directories.
2. Edit parameter settings in your adapter configuration file:

   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."

   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 168.
3. To update an adapter so that the Solaris operating system can recognize all the LUNs, perform the following steps:
   a. Type `cfgadm -c configure` *`adapter_number`*.
   b. Type `devfsadm` to scan all attached devices.
   c. Type `format` to display all LUNs that were discovered by the Solaris operating system.

Parameter settings for Sun-Emulex adapters on a Sun host

You can use these recommended settings for Sun-Emulex SG-XPCI1/2FC-EM*x*, and SG-XPCIe1/2FC-EM*x* adapters on a Sun host.

Table 18 describes configuration parameters and recommended settings for Sun-Emulex adapters.

Table 18. 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-alpa=0x01;  assigns ALPA 0x01 to adapter 0`
	`emlxs1-assign-alpa=0x02;  assigns ALPA 0x02 to adapter 1`
	`emlxs2-assign-alpa=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.
cr-delay	**0** (Default. The coalesce response feature is disabled.) 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.

Table 18. Recommended configuration settings for Sun-Emulex adapters on a Sun host (continued)

Parameters	Recommended settings
link-speed	**0** (Default) Specifies the link speed to be used for initializing the Fibre Channel connection: 0: Select automatically 1: 1 Gigabaud 2: 2 Gigabauds 4: 4 Gigabauds
num-nodes	**512** (Default) Specifies the number of Fibre Channel nodes (N_Ports). You can specify a value of 2 - 512.
pm-support	**1** (Default) Enables or disables power management support in the driver: 0: Support is disabled. 1: Support is enabled.
topology	**0** (Default, unless you want to run as an N_Port or an NL_Port.) Specifies the link topology to be used for initializing the Fibre Channel connection: 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	**1** (Default) Enables or disables IP networking support in the driver: 0: Support is disabled. 1: Support is enabled.
Message parameters	
console-errors	**0** (Default) Specifies the verbose level for driver errors to the console: 0: Sets the minimum level. 0xffffffff: Sets the maximum level.
console-notices	**0** (Default) Specifies the verbose level for driver notices to the console: 0: Sets the minimum level. 0xffffffff: Sets the maximum level.
console-warnings	**0** (Default) Specifies the verbose level for driver warnings to the console: 0: Sets the minimum level. 0xffffffff: Sets the maximum level.

Table 18. Recommended configuration settings for Sun-Emulex adapters on a Sun host (continued)

Parameters	Recommended settings
log-errors	**0xffffffff** (Default) This parameter sets the verbose level for driver errors to the system log file: 0: Sets the minimum level. 0xffffffff: Sets the maximum level. To configure the types of events to be logged, see "System log file parameters."
log-notices	**0xffffffff** (Default) Specifies the verbose level for driver notices to the system log file: 0: Sets the minimum level. 0xffffffff: Sets the maximum level. To configure the types of events to be logged, see "System log file parameters."
log-warnings	**0xffffffff** (Default) Specifies the verbose level for driver warnings to the system log file: 0: Sets the minimum level. 0xffffffff: Sets the maximum level. To configure the types of events to be logged, see "System log file parameters.".

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 19 lists the system log parameters and settings you can use to specify the types of messages to be logged.

Table 19. 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
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

Table 19. System log parameters for configuring Sun-Emulex adapter messages (continued)

Event Types	Parameters	Settings
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 Sun-QLogic adapters on a Sun host

You can use these recommended configuration settings for Sun-QLogic SG-XPCI1/2FC-QF*x* and SG-XPCIe1/2FC-QF*x* 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 DS8000 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	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. Specifies the connection mode that the driver firmware uses: 0: Loop only 1: Point-to-point only 2: Loop preferred, otherwise point-to-point (default)
enable-adapter-hard-loop-ID	**0** (Default) Enables or disables the setting of the hard loop ID for adapters on the Fibre Channel bus: 0: Enabled 1: Disabled The default value is recommended and might be necessary on some bus configurations where devices fail to appear.
enable-FCP-2-error-recovery	**1** (Default) Enables or disables the Fibre Channel FCP 2 error recovery support: 0: Disabled 1: Enabled

Table 20. Recommended configuration settings for Sun-QLogic adapters on a Sun host (continued)

Parameters	Recommended settings
enable-link-down-error	**1** (Default) Enables or disables driver error reporting during a link-down condition: 0: Disabled 1: Enabled
enable-LIP-full-login-on-bus-reset	**1** (Default) Enables or disables the adapter to issue a LIP full login reset during Fibre Channel reset: 0: Disabled 1: Enabled
enable-LIP-reset-on-bus-reset	**0** (Default) Enables or disables the adapter to issue a LIP reset during Fibre Channel reset: 0: Disabled 1: Enabled
enable-target-reset-on-bus-reset	**0** (Default) Enables or disables the adapter to issue a target reset during Fibre Channel reset: 0: Disabled 1: Enabled The default value is recommended.
execution-throttle	**32** (Default) Specifies the maximum number of commands sent per LUN by the ISP2*xxx* firmware. You can specify a value of 1 - 65535. **Note:** Exceeding device capabilities causes unnecessary command retries that impact performance.
extended-logging	**0** (Default) Enables or disables the logging of driver-detected events that occur in the driver or Fibre Channel bus: 0: Disabled 1: Enabled Events are usually logged in the `/var/adm/messages` file.
fc-data-rate	**2** (Default) Specifies the Fibre Channel data rate which the driver uses: 0: 1 gigabit/second 1: 2 gigabits/second 2: Automatically negotiate 3: 4 gigabits/second

Table 20. Recommended configuration settings for Sun-QLogic adapters on a Sun host (continued)

Parameters	Recommended settings

firmware-dump-flags	**0** (Default)
	
	Specifies whether to force a firmware dump when one or more flag conditions occur:
	0: Disabled
	1: Enabled
	
	Flag conditions:
	`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`

link-down-timeout	**0** (Default) 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.
login-retry	**4** (Default) 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.
login-timeout	**3** (Default) 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. **Note:** Large values can cause long delays during initialization and device reconfiguration.
max-frame-length	**2048** (Default) Specifies the frame payload length, in bytes, that is used by the ISP2*xxx* firmware. Other possible values are 1024 and 2048.
pci-max-read-request	**2048** (Default) 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. **Note:** If this value is not specified or is not equal to 128, 256, 512, 1024, 2048 or 4096, the ISP2*xxx* defaults to the value specified by the system.
port-down-retry-count	**8** (Default) Specifies the number of command retries when devices are not responding on the Fibre Channel bus. You can specify a value of 0 - 255. **Note:** Large values can cause long delays for failover software to detect a failing device.

Table 20. Recommended configuration settings for Sun-QLogic adapters on a Sun host (continued)

Parameters	Recommended settings
port-down-retry-delay	**2** (Default) Specifies the number of seconds to delay between port-down retries. You can specify a value of 0 - 255. **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.
queue-full-retry-count	**16** (Default) Specifies the number of times to retry an SCSI queue-full error. You can specify a value of 0 - 255.
queue-full-retry-delay	**2** (Default) 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.
reset-delay	**5** (Default) 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.

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
```

Tip: If you are using Solaris 8 or 9, you must install the SAN software package before you can display information on all the parameters that are associated with the SSD driver. For more information, see "Downloading and installing a host bus adapter driver" on page 2.

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.

About this task

To set the system parameters, perform the following steps:

Procedure

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 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 connected LUN. Determine the value by using the following formula: `256 ÷ (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, add the following line to the /etc/system file:

   ```
   set sd:sd_max_throttle=8
   ```

   Notes:

   a. In general, an sd_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.
   b. Theoretically, all LUNs on a port must be simultaneously busy to reach the maximum throttle setting that is specified on the sd_max_throttle parameter times the number of LUNs. You can tune the sd_max_throttle setting higher than the formula dictates; however, this might result in an unexpected workload spike that can cause queue overrun.

   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.

About this task

To set the system parameters, perform the following steps:

Procedure

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 ÷ (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, add the following line to the /etc/system file:

   ```
   set sd:sd_max_throttle=8
   ```

   Notes:

   a. In general, an sd_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.

   b. Theoretically, all LUNs on a port must be simultaneously busy to reach the maximum throttle setting that is specified on the sd_max_throttle parameter times the number of LUNs. You can tune the sd_max_throttle setting higher than the formula dictates; however, this might result in an unexpected workload spike that can cause queue overrun.

   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 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 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.

About this task

To set the system parameters, perform the following steps:

Procedure

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 ÷ (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, add the following line to the /etc/system file:

   ```
   set ssd:ssd_max_throttle=8
   ```

   Notes:

a. 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.
b. 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.

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.

About this task

Note: If you have installed multiple host adapters, you will see more than one WWPN.

To locate the WWPN, perform the following steps:

Procedure

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 WWPN.
   - AMCC adapters:
     - For the 2Gb AMCC adapter, search for `AMCCc146xN: Fibre Channel WWNN`, where *N* is the adapter number (0, 1, and so on). You can find the WWPN on the same line following the WWNN.
     - For the AMCC PCI adapter, search for `fca-pci`*x*`: Fibre Channel WWNN`, where *x* is the adapter number (0, 1, and so on). You can find the WWPN on the same line following the WWNN.
     - For the AMCC SBUS adapter, search for `fcaw`*x*`: Fibre Channel WWNN`, where *x* is the adapter number (0, 1, and so on). You can find the WWPN on the same line immediately after the WWNN.
   - For the Emulex adapter, search for `lpfc`*x*`: Fibre Channel WWNN`, where *x* is the adapter number (0, 1, and so on).
   - For the QLogic QLA23*xx* adapter, search for `qla23200-hba`*x*`-adapter-port-name` where *x* is the adapter number (0, 1, and so on).
   - For the Sun-Emulex adapter, search for `emlxs`*x*`: Fibre Channel WWNN`, where *x* is the adapter number (0, 1, and so on).
   - For the Sun-QLogic adapter, search for `qlc`*x*`: Fibre Channel WWNN`, 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 DS8000 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 DS Storage Manager, 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 website at 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 DS8000.

- 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 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 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.

About this task

To configure the STMS host settings, perform the following steps:

Procedure

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 ÷ (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, add the following line to the /etc/system file:

   ```
   set ssd:ssd_max_throttle=8
   ```

   Notes:

   a. 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.

b. 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.

4. To enable Sun MPxIO multipathing for an STMS host that is attached to a DS8000:
   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[1]
auto-failback	enable
[1] For a symmetric storage array, you must specify round-robin so that the driver can balance the I/O load between the two paths.	

   c. Add the following lines to the end of the **scsi_vhci.conf** file:

```
device-type-scsi-options-list =
"IBM     2107900", "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 topic describes how to attach a storage unit to a Sun host using Sun Cluster.

See the Sun website at www.sun.com for the most current documentation.

Notes:

1. Sun Cluster software requires MPXIO as the multipath driver in all configurations regardless of the version.
2. For a two-node SunCluster 3.0 configuration, Sun and third-party HBAs are supported if the Sun Fibre Channel and MPXIO drivers are available to the operating system and HBA models that are in the configuration. The Sun Fibre Channel and MPXIO drivers are part of the Sun SAN installation package which can be downloaded from the Sun website. For more information, see "Downloading and installing a host bus adapter driver" on page 2.
3. For the most current Sun Cluster interoperability matrix for IBM storage products, go to Oracle's Solaris Cluster Storage Partner Program.

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"
DS8000 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
```

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 one 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
```

Notices

The information provided by this media supports the products and services described with consideration for the conditions described herein.

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 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 PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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.

Trademarks

This topic lists trademarks that appear in this information.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web by going to the Copyright and trademark information website.

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 and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

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.

Other company, product, and service names may be trademarks or service marks of others.

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission statement

This explains the Federal Communications Commission's (FCC) statement.

This equipment has been tested and found to comply 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 in order 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 user's 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 Electromagnetic Compatibility Directive

This product is in conformity with the protection requirements of European Union (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 non-recommended modification of the product, including the fitting of non-IBM option cards.

Attention: This is an EN 55022 Class A product. In a domestic environment this product might cause radio interference in which case the user might be required to take adequate measures.

Responsible Manufacturer:

International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
914-499-1900

European community contact:

IBM Technical Regulations, Department M456
IBM-Allee 1, 71137 Ehningen, Germany
Tel: +49 7032 15-2937
E-mail: tjahn@de.ibm.com

Germany Electromagnetic compatibility directive

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2004/108/EG 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 Mabnahmen 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 2004/108/EG in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC EG Richtlinie 2004/108/EG) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:

```
International Business Machines Corp.
New Orchard Road
Armonk,New York 10504
Tel: 914-499-1900
```

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:

IBM Deutschland
Technical Regulations, Department M456

IBM-Allee 1, 71137 Ehningen, Germany
Tel: +49 7032 15-2937
e-mail: tjahn@de.ibm.com

Generelle Informationen:

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、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。 VCCI-A

Translation:

This is a Class A product based on the standard of the VCCI Council. If this equipment is used in a domestic environment, radio interference may occur, in which case, the user may be required to take corrective actions.

Japanese Electronics and Information Technology Industries Association (JEITA) statement

Japanese Electronics and Information Technology Industries Association (JEITA) Confirmed Harmonics Guideline (products less than or equal to 20 A per phase)

高調波ガイドライン適合品

Japanese Electronics and Information Technology Industries Association (JEITA) Confirmed Harmonics Guideline with Modifications (products greater than 20 A per phase)

高調波ガイドライン準用品

Korea Communications Commission (KCC) Electromagnetic Interference (EMI) Statement

이 기기는 업무용(A급)으로 전자파적합기기로서 판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의 지역에서 사용하는 것을 목적으로 합니다.

f2c01252

Russia Electromagnetic Interference (EMI) Class A Statement

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

rusemi

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。

taiemi

Taiwan contact information

This topic contains the product service contact information for Taiwan.

IBM Taiwan Product Service Contact Information:
IBM Taiwan Corporation
3F, No 7, Song Ren Rd., Taipei Taiwan
Tel: 0800-016-888

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

f2c00790

Index

Printed in USA

GC27-2298-02