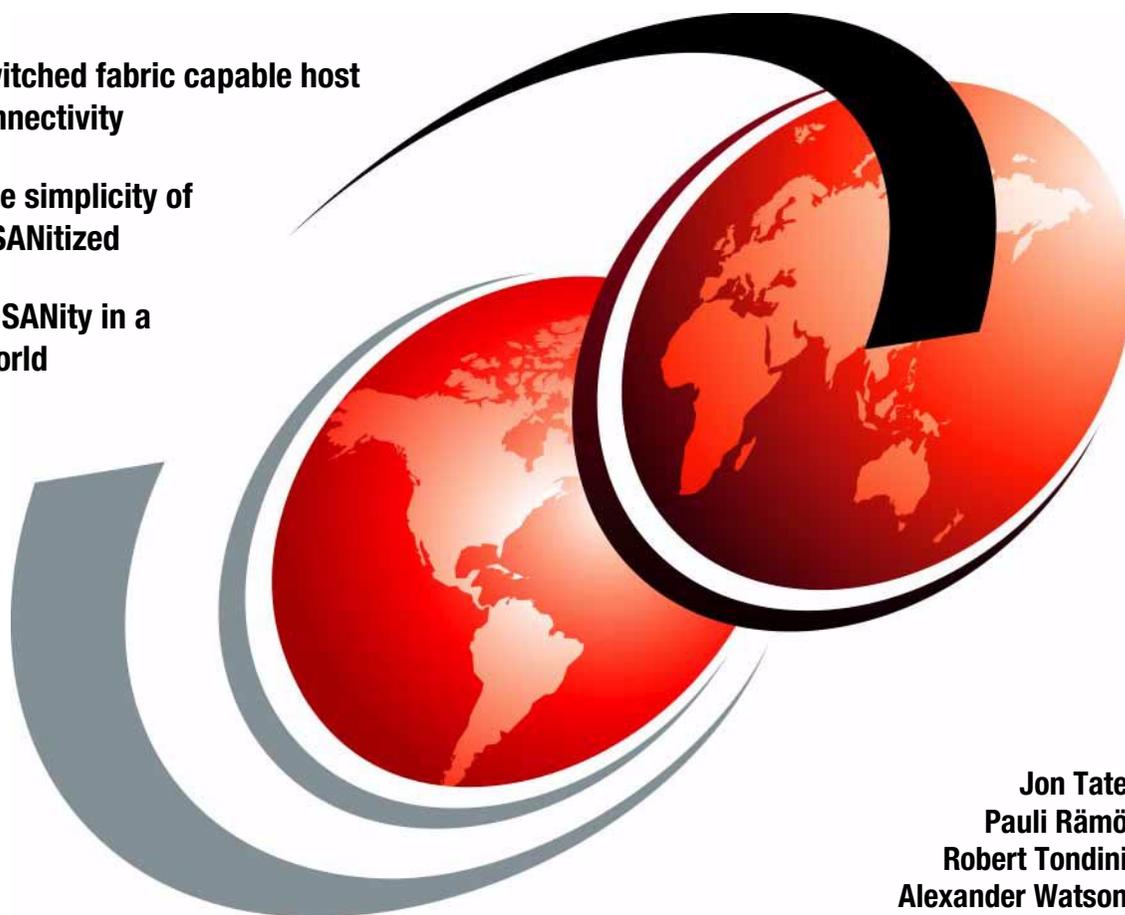


Introducing Hosts to the SAN Fabric

Ready a switched fabric capable host
for SAN connectivity

Discover the simplicity of
becoming SANitized

Verify your SANity in a
complex world



Jon Tate
Pauli Rämö
Robert Tondini
Alexander Watson



International Technical Support Organization

Introducing Hosts to the SAN Fabric

February 2002

Take Note! Before using this information and the product it supports, be sure to read the general information in “Special notices” on page xiii.

First Edition (February 2002)

This edition applies to those IBM SAN software and hardware products currently in the IBM portfolio and described herein.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. QXXE Building 80-E2
650 Harry Road
San Jose, California 95120-6099

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 2002. All rights reserved.

Note to U.S Government Users – Documentation related to restricted rights – Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	vii
Tables	ix
Examples	xi
Special notices	xiii
IBM trademarks	xiv
Preface	xv
The team that wrote this redbook	xv
Notice	xviii
Comments welcome	xviii
Chapter 1. Windows Fibre Channel implementation	1
1.1 Fibre Channel host bus adapter introduction	2
1.2 Windows supported host bus adapters	2
1.2.1 QLogic Fibre Channel host bus adapter cards	3
1.2.2 Emulex Fibre Channel host bus adapter cards	4
1.2.3 JNI Fibre Channel host bus adapter cards	6
1.3 Supported SAN fabric components	6
1.3.1 General information	6
1.3.2 The IBM 2109 family	7
1.3.3 INRANGE	7
1.3.4 McDATA	8
1.4 Storage	9
1.5 FC host bus adapter card installation procedure	11
1.5.1 Preparing to install the Fibre Channel HBAs	11
1.5.2 Installing QLogic Fibre Channel HBAs family on Windows 2000	11
1.5.3 Installing QLogic Fibre Channel HBAs family on Windows NT	12
1.5.4 QLogic configuration utility	13
1.5.5 Installing Emulex Fibre Channel HBAs in Windows 2000	16
1.5.6 Installing Emulex in Windows NT	17
1.5.7 Emulex configuration utility	18
1.5.8 Installing JNI Fibre Channel host bus adapters in Windows 2000	21
1.5.9 Installing JNI Fibre Channel host bus adapters in Windows NT	22
1.5.10 JNI configuration utility	22
1.6 Installing the storage system	23

Chapter 2. eServer pSeries Fibre Channel implementation	33
2.1 Introduction	34
2.2 Fibre Channel host bus adapters	34
2.2.1 IBM FC HBAs	34
2.2.2 Other FC HBAs	35
2.3 SAN fabrics	36
2.3.1 General information.	36
2.3.2 The IBM 2109 family	37
2.3.3 INRANGE	38
2.3.4 McDATA	38
2.4 Storage	40
2.5 Installation procedures	40
2.5.1 Preparations to install the FC HBAs	40
2.5.2 Installing the FC HBAs	42
2.5.3 Installing the storage system.	53
Chapter 3. Sun Fibre Channel implementation	59
3.1 Overview of the Sun system	60
3.2 Sun supported Fibre Channel host adapters.	62
3.2.1 Tuning HBA configuration files	65
3.3 Sun supported Fibre Channel switches and hubs	69
3.3.1 IBM 2109.	69
3.3.2 INRANGE FC/9000.	71
3.3.3 McDATA	72
3.3.4 Fibre Channel switch installation preparation	74
3.4 Supported storage on the Sun fabric.	74
3.5 Installation planning.	75
3.5.1 Cabling requirements and considerations.	76
3.6 Installation procedure	76
3.6.1 Check that the system is at a proper level	77
3.6.2 Installing the FC HBA adapter card.	77
3.6.3 Switch setup	77
3.6.4 Storage setup	80
3.6.5 Downloading the Fibre Channel adapter drivers.	81
3.6.6 Installing the Fibre Channel adapter driver(s).	83
3.6.7 Configuring the host device and adapter drivers.	86
3.7 Device functionality checks	90
3.8 Specific limitations.	91
Chapter 4. HP 9000 Fibre Channel implementation	93
4.1 Introduction	94
4.2 Fibre Channel host bus adapters	94
4.2.1 HP FC HBAs	94

4.2.2 Other FC HBAs	95
4.3 SAN fabrics	95
4.3.1 General information.	96
4.3.2 The IBM 2109 family	97
4.3.3 INRANGE	97
4.3.4 McDATA	98
4.4 Storage	99
4.5 Installation procedures	100
4.5.1 Preparing to install the FC HBA	100
4.5.2 Installing the FC HBA	101
4.5.3 Installing the storage system.	103
Chapter 5. NUMA-Q Fibre Channel implementation	107
5.1 Overview of the NUMA-Q system	108
5.2 Fibre Channel host bus adapters	109
5.2.1 Host adapter placement	111
5.2.2 Host adapter installation preparation	112
5.3 Fibre Channel switch support	112
5.3.1 Fibre Channel switch installation preparation	113
5.4 Supported storage on the NUMA-Q fabric	114
5.5 Installation planning.	116
5.6 Installation procedure	117
5.6.1 Verifying the correct software is installed	117
5.6.2 FC HBA setup	118
5.6.3 Switch setup	120
5.6.4 Storage setup	122
5.7 Device functionality checks	125
5.8 Specific limitations.	126
Chapter 6. Compaq AlphaServer Fibre Channel Implementation.	127
6.1 Introduction	128
6.2 Fibre Channel host bus adapters	128
6.2.1 Compaq FC HBA	128
6.3 Supported SAN switches.	129
6.3.1 General information.	130
6.3.2 IBM 2109-S08/S16	131
6.4 Storage	131
6.5 Installation procedures	132
6.5.1 Preparing to install the FC HBA	132
6.5.2 Installing the FC HBA	133
6.6 Installing the storage system.	136
Chapter 7. Linux on i386 platform Fibre Channel implementation	139
7.1 Introduction	140

7.2 Fibre Channel host adapters	141
7.2.1 QLogic	141
7.2.2 Emulex	142
7.2.3 JNI	142
7.3 SAN fabric	142
7.3.1 General information	143
7.3.2 IBM 2109 family	144
7.3.3 INRANGE	144
7.3.4 McDATA	145
7.4 Storage	146
7.5 Installation procedures	146
7.5.1 Installing a QLogic FC HBA	147
7.5.2 Configuring storage for the QLogic FC HBA	151
7.5.3 Installing an Emulex FC HBA	155
7.5.4 Configuring storage for the Emulex FC HBA	160
Glossary	163
Related publications	175
IBM Redbooks	175
Other resources	175
Referenced Web sites	178
How to get IBM Redbooks	180
IBM Redbooks collections	180
Index	181

Figures

1-1	QLogic SANblade 2200 Series	3
1-2	QLogic SANblade 2300 Series	4
1-3	Emulex LP9000	5
1-4	QLview for Fibre utility — Set Up Connection screen	13
1-5	QLview for Fibre utility — Main Window	14
1-6	QLview for Fibre utility	15
1-7	QLview for Fibre utility — adapter information	16
1-8	Emulex LP8000 Fibre Channel HBAs — Status LEDs	18
1-9	Light Pulse utility — main window	19
1-10	LightPulse Utility — Category list	19
1-11	LightPulse Utility — Firmware download	20
1-12	IBM 2109 Switch — configuration main menu	25
1-13	IBM 2109-S16 Switch Management Application	25
1-14	IBM 2109-S16 — Port Information	26
1-15	Write Signature and Upgrade Disk Wizard	27
1-16	Select Disk	28
1-17	Completing the Write Signature task	29
1-18	Disk Management — disk 1 unallocated and unformatted	30
1-19	Disk Administrator change warning	30
1-20	New device recognized	31
1-21	Creating partitions	31
1-22	After creation of partitions	32
1-23	Commit changes	32
2-1	IBM eServer pSeries Support	41
2-2	Install additional device software — select install medium	43
2-3	Install additional device software — initial settings	43
2-4	Install additional device software — choose software package	44
2-5	Install additional device software — installation settings	45
2-6	Install additional device software — installation in progress	45
2-7	Install additional device software — mount additional volume	46
2-8	Install additional device software — results	46
2-9	Update device software — install medium	47
2-10	Update device software — installation settings	48
2-11	Update device software — installation in progress	48
2-12	Update device software — results	49
2-13	Install storage drivers — select install medium	54
2-14	Install storage drivers — installation in progress	54
2-15	Install storage drivers — results	55

2-16	Update storage drivers — select install medium	56
2-17	Update storage drivers — installation settings	56
2-18	Update storage drivers — installation in progress	57
2-19	Update storage software — results	57
3-1	Emulex LP8000 adapter	63
3-2	JNI FC64-1063 adapter	64
3-3	QLogic QLA2200F adapter	65
3-4	IBM 2109-S08 (top) and S16	70
3-5	INRANGE FC/9000 switch	71
3-6	McDATA cabinet (left), and 6064 rack mountable switch	73
5-1	Emulex LP7000E Fibre Channel Adapter	110
5-2	Emulex LP6000 Fibre Channel Adapter	111
5-3	IBM 2109-S08 (top) and S16	113
6-1	Compaq StorageWorks 64-Bit/33MHz PCI-to-FC HBA	129
6-2	64-Bit PCI Fibre Channel Host Adapter	134

Tables

1-1	Windows supported Fibre Channel host bus adapters	2
1-2	Common Fibre Channel cabling configurations	6
1-3	The IBM 2109 family	7
1-4	The INRANGE switches.	8
1-5	The McDATA switches.	8
1-6	IBM Fibre Channel attached disk storage systems	10
1-7	IBM Fibre Channel attached tape storage systems	10
2-1	IBM FC HBAs for IBM eServer pSeries and RS/6000	34
2-2	Other FC HBAs available for IBM eServer pSeries and RS/6000	35
2-3	Common Fibre Channel cabling configurations	36
2-4	The IBM 2109 family	37
2-5	The INRANGE switches.	38
2-6	The McDATA switches.	39
2-7	IBM Fibre Channel attached disk storage systems	40
2-8	IBM Fibre Channel attached tape storage systems	40
3-1	Solaris 2.6, 7, and 8 minimum patch levels for Fibre Channel.	61
3-2	Solaris kernel parameter settings required	65
3-3	Recommended JNI HBA settings	66
3-4	JNI FCE-6460 settings.	66
3-5	Recommended Emulex LP8000 HBA settings.	67
3-6	Recommended QLA2200F HBA settings	68
3-7	The IBM 2109 family	69
3-8	INRANGE directors	71
3-9	McDATA components	72
3-10	IBM Fibre Channel attached disk storage systems	75
3-11	IBM Fibre Channel attached tape storage systems	75
3-12	Common Fibre Channel cabling configurations	76
4-1	Current HP FC HBAs with fabric support.	94
4-2	Other FC HBAs available for HP 9000 systems	95
4-3	Common Fibre Channel cabling configurations	96
4-4	The IBM 2109 family	97
4-5	The INRANGE switches.	98
4-6	The McDATA switches.	98
4-7	IBM Fibre Channel attached disk storage systems	99
4-8	IBM Fibre Channel attached tape storage systems	100
4-9	Required patches for HP-UX versions	100
5-1	Supported FC HBAs	109
5-2	The NUMA-Q supported switch family.	112

5-3	IBM Fibre Channel attached disk storage systems	114
5-4	IBM NUMA-Q RPQ Fibre Channel attached disk storage systems . . .	115
5-5	IBM NUMA-Q standard Fibre Channel attached tape storage systems	115
5-6	IBM NUMA-Q RPQ Fibre Channel attached tape storage systems . . .	116
5-7	Ready made fibre optic duplex cables	116
5-8	Common Fibre Channel cabling configurations	117
6-1	Common Fibre Channel cabling configurations	130
6-2	The IBM 2109 family	131
6-3	IBM Fibre Channel attached disk storage systems	131
6-4	64-Bit PCI Fibre Channel HBA description	134
7-1	FC HBAs supported in Linux	141
7-2	Common Fibre Channel cabling configurations	143
7-3	The IBM 2109 family	144
7-4	The INRANGE switches.	144
7-5	The McDATA switches.	145
7-6	IBM Fibre Channel attached disk storage systems	146

Examples

1-1	Switch connectivity check	24
2-1	Checking the AIX maintenance level	42
2-2	Checking the Fibre Channel adapter	50
2-3	Checking the firmware level	50
2-4	Checking the protocol	51
2-5	Checking the FC HBA connectivity to the switch	52
2-6	Creating and testing a filesystem	58
4-1	Checking the FC HBA	101
4-2	Checking the FC HBA parameters	102
4-3	Checking the FC HBA connectivity to the switch	103
4-4	Creating volume group and filesystem on the new disk	105
6-1	Result of set mode diag and wwidmgr -show adapter commands	135
6-2	Formatting and initializing FC HBA	136
6-3	Switch connectivity check	137
6-4	Storage connectivity check	138
6-5	Decimal volume number for an ESS volume	138
6-6	Verifying disk attachment by using hwmgr command	138
7-1	Checking the kernel messages from loading of QLogic driver	149
7-2	Checking the QLogic adapter via /proc interface	150
7-3	Checking QLogic adapter switch connectivity	150
7-4	Checking the QLogic adapter again via /proc interface	152
7-5	Checking the kernel messages	152
7-6	Testing a Fibre Channel disk with QLogic adapter	153
7-7	Checking the Emulex adapter with the lputil command	158
7-8	Checking the Emulex adapter switch connectivity	159
7-9	Testing a Fibre Channel disk with Emulex adapter	161

Special notices

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

IBM trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

e (logo)® 	FICON™	RS/6000®
Redbooks (logo)™ 	IBM®	S/390®
AIX®	Magstar®	SANergy™
AIX 5L™	Netfinity®	SP™
DB2®	Notes®	SP2®
DYNIX®	NUMA-Q®	StorWatch™
DYNIX/ptx®	OS/2®	System/390®
Enterprise Storage	OS/390®	Tivoli®
Server™	Perform™	TotalStorage™
ESCON®	pSeries™	Wave®
	Redbooks™	

Other company trademarks

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Preface

In this IBM Redbook, we aim to be short, sharp, and to the point — and provide an answer. There is no marketing, no sales pitch, no extraneous information, other than what is needed to get the job done. That job is making hosts Storage Area Network (SAN) ready.

To this end, we show how to make hosts SAN fabric ready and connect them to the fabric with minimum fuss. We cover all SAN switched fabric capable hosts.

For background reading on SAN, we suggest these books as relevant:

- ▶ *IBM SAN Survival Guide*, SG24-6143
- ▶ *IBM SAN Survival Guide Featuring the IBM 2109*, SG24-6127
- ▶ *IBM SAN Survival Guide Featuring the McDATA Portfolio*, SG24-6149
- ▶ *IBM SAN Survival Guide Featuring the INRANGE Portfolio*, SG24-6150
- ▶ *Designing an IBM Storage Area Network*, SG24-5758
- ▶ *Implementing an Open IBM SAN*, SG24-6116
- ▶ *Introduction to Storage Area Network, SAN*, SG24-5470

With the vast combination of SAN components, system levels, hardware and software, it is not possible to cover each and every environment that is likely to be encountered in the field. However, we do intend to cover those elements that are common, and that any reader with an intermediate grasp of SAN and host concepts, should be able to apply.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.



The team (from left to right): Jon, Robert, Al, Pauli

Jon Tate is a Project Leader for SAN TotalStorage Solutions at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing level 2 support for IBM storage products. Jon has 16 years of experience in storage software and management, services, and support, and is an IBM SAN Certified Specialist.

Pauli Rämö is an IT Specialist in IBM Finland. He has 9 years of experience in RS/6000, AIX, and Linux. His areas of expertise include RS/6000 and RS/6000 SP systems, AIX, HACMP, open systems storage solutions, and SAP R/3 Basis. He is an IBM Certified Advanced Technical Expert - RS/6000 AIX V4 and a SAP Certified Technical Consultant UNIX/Oracle and UNIX/DB2 UDB.

Robert Tondini is a Storage Specialist in IBM Croatia. He has 7 years of experience in S/390 platform, initially as a VM/VSE system programmer and later on migrating to the OS/390 environment as a DB2 system administrator. He has been employed at IBM for 1 year. His areas of expertise include IBM disk and tape storage subsystem solutions.

Alexander Watson is a System Engineer with Remotely Delivered Services (RDS), Beaverton Support Center (BSC), Beaverton, Oregon, USA. He has 28 years of experience in the computer service field, spanning from mainframe to personal systems. He currently specializes in the IBM NUMA-Q, and pSeries Open Systems SAN and storage networking support. His areas of expertise include hardware support, I/O device networking, mixed open systems SAN environments, RPQ storage, and ESS support on open systems platforms.

Thanks to the following people for their contributions to this project:

Will Carney
Emma Jacobs
Yvonne Lyon
Barry Mellish
Deanna Polm
Sokkieng Wang
International Technical Support Organization, San Jose Center

Christopher Dace
IBM SSG

Scott Drummond
IBM SSG

Glenda Fuller
IBM SSG

Alejandro Halili
IBM SSG

Richard Heffel
IBM SSG

Michael Janini
IBM SSG

Notice

This publication is intended to help systems and storage administrators install IBM SAN portfolio equipment. The information in this publication is not intended as the specification of any programming interfaces that are provided by any of the SAN hardware and software components contained herein. See the PUBLICATIONS section of the IBM Programming Announcement for the SAN hardware and software components contained herein for more information about what publications are considered to be product documentation.

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:
ibm.com/redbooks
- ▶ Send your comments in an Internet note to:
redbook@us.ibm.com
- ▶ Mail your comments to the address on page ii.



Windows Fibre Channel implementation

In this chapter we describe the steps involved in implementing Fibre Channel host bus adapter cards in an Intel based server running Windows NT or Windows 2000 operating system and its connection to the SAN.

The following topics are discussed:

- ▶ “Fibre Channel host bus adapter introduction” on page 2
- ▶ “Windows supported host bus adapters” on page 2
- ▶ “Supported SAN fabric components” on page 6
- ▶ “Storage” on page 9
- ▶ “FC host bus adapter card installation procedure” on page 11
- ▶ “Installing the storage system” on page 23

1.1 Fibre Channel host bus adapter introduction

In the topics that follow, we introduce Windows NT and Windows 2000 supported Fibre Channel host bus adapter cards. These are manufactured by QLogic, Emulex, and JNI. We also cover their implementation, information with respect to the latest drivers and firmwares levels.

We also detail supported SAN switches and the considerations which should be taken into account in order to establish the connection with Intel based servers running Windows NT or 2000.

For more general information or technical information on employment and application of these adapters, refer to the appropriate Fibre Channel adapter manufacturer.

1.2 Windows supported host bus adapters

Windows NT and 2000 environments support many Fibre Channel host bus adapters. There are many reasons as to why one would be chosen over another, depending on price, supported topologies, brand loyalty, existing adapters already installed, and so on. We do not recommend one type or another.

In Table 1-1 we show a list of the latest QLogic, Emulex, and JNI models which are compatible and supported by Windows NT and 2000.

Table 1-1 Windows supported Fibre Channel host bus adapters

Adapter Manufacturer	Adapter Type	No of ports	Link speed	Connector type
QLogic	QLA 2200F	1	1 Gb/s	SC
QLogic	QLA 2202F	2	1 Gb/s	SC
QLogic	QLA 2204F	4	1 Gb/s	LC
QLogic	QLA 2300F	1	2 Gb/s	LC
QLogic	QLA 2302F	2	2 Gb/s	LC
Emulex	LP7000E	1	1 Gb/s	SC
Emulex	LP850	1	1 Gb/s	SC
Emulex	LP8000	1	1 Gb/s	SC
Emulex	LP8000DC	2	1 Gb/s	SC

Adapter Manufacturer	Adapter Type	No of ports	Link speed	Connector type
Emulex	LP9002DC	2	1-2 Gb/s	LC
Emulex	LP9002L	1	1-2 Gb/s	LC
Emulex	LP9402DC	2	1-2 Gb/s	LC
Emulex	LP952L	1	1-2 Gb/s	LC
JNI	FCE-3210	1	1 Gb/s	SC
JNI	FCE-6410	1	1 Gb/s	SC
JNI	FCE2-6410	2	1 Gb/s	SC
JNI	FCE2-6412	2	1 Gb/s	SC

1.2.1 QLogic Fibre Channel host bus adapter cards

There is a wide portfolio of QLogic Fibre Channel host bus adapter cards supported by Windows NT and 2000. In this section we cover the SANblade 2200 (shown in Figure 1-1) and the SANblade 2300 family (shown in Figure 1-2).

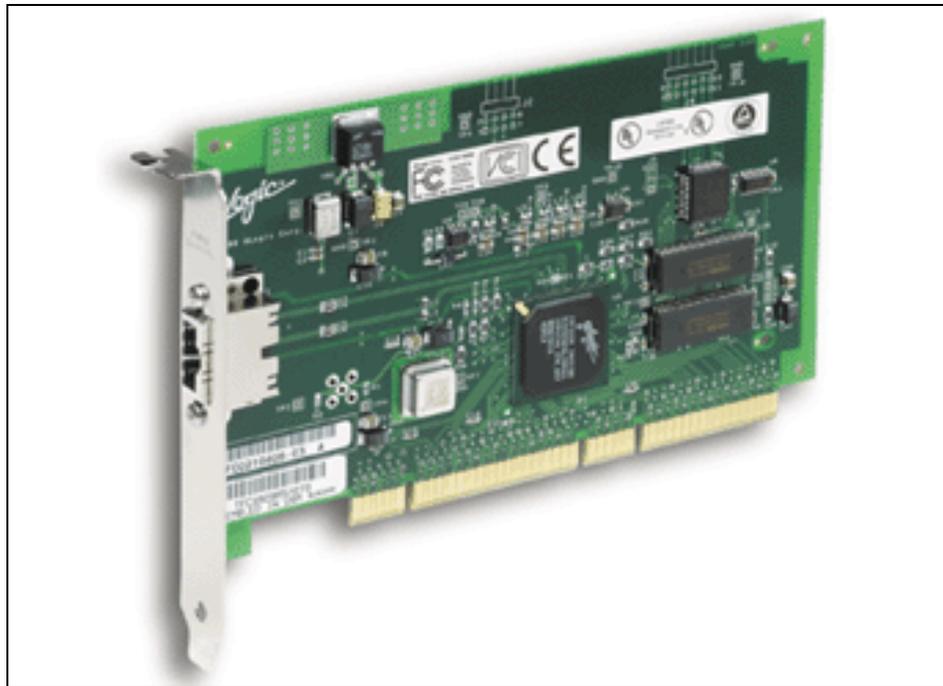


Figure 1-1 QLogic SANblade 2200 Series

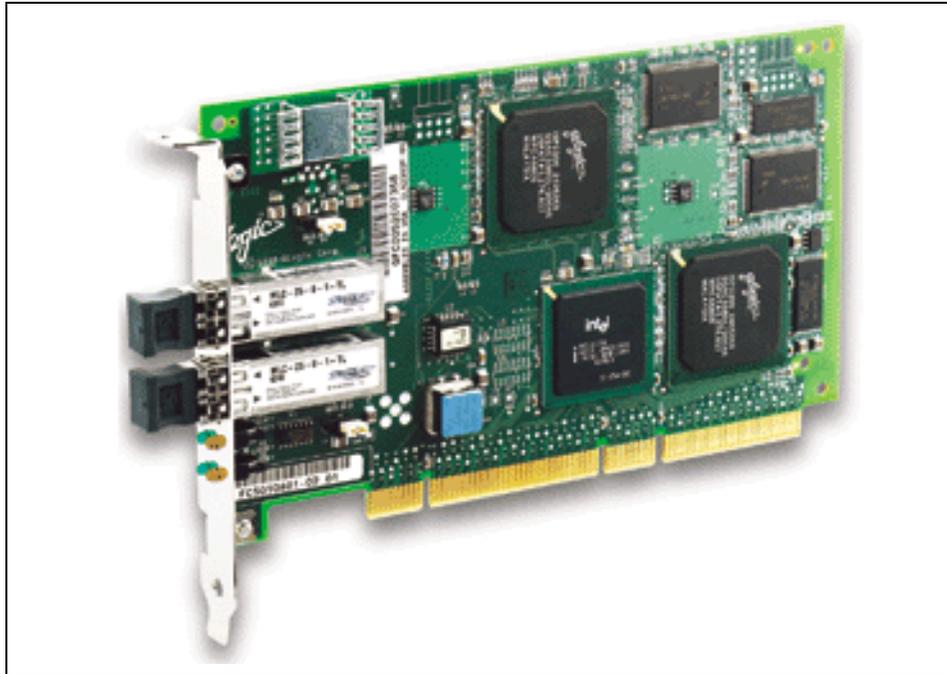


Figure 1-2 QLogic SANblade 2300 Series

The following Web site offers more technical information, including the latest driver version, and whether it is Microsoft certified for use as a storage adapter and cluster device in the Windows NT and Windows 2000 environments:

<http://www.qlogic.com>

These are some of the supported features that are common for all QLogic FC host bus adapter cards mentioned in Table 1-1 on page 2:

- ▶ Support of F_Port, FL_Port, and N_Port configurations
- ▶ Extended LUN support
- ▶ IP support
- ▶ FC Tape support
- ▶ Management Server interface support

1.2.2 Emulex Fibre Channel host bus adapter cards

Two of the Windows NT and 2000 supported FC host bus adapter cards that belong to the Emulex family are the LP8000 and LP9000.

We show the LP9000 in Figure 1-3.



Figure 1-3 Emulex LP9000

Both are designed for large scale applications. The basic technical specifications for the Windows supported models can be found in Table 1-1 on page 2.

For more detail, we suggest that you visit the Emulex Web site in order to download the latest firmware and driver level for each adapter: The URL is:

<http://www.emulex.com>

Following are the prerequisites related to the operating system level for the Emulex FC adapters mentioned:

- ▶ Windows NT 4.0 Service Pack 3 (service Pack 6 recommended)
- ▶ Windows 2000 Service Pack 2 (recommended but not necessary)

These are the basic features in common for all the Emulex Fibre Channel HBAs listed in Table 1-1 on page 2:

- ▶ Point-to-Point, Arbitrated-Loop and Switched Fabric topologies are supported
- ▶ FC tape support
- ▶ Concurrent use of multiple protocols (SCSI and IP)

1.2.3 JNI Fibre Channel host bus adapter cards

JNI offers four FC host cards supported in Windows NT and 2000 environments. The minimum system requirement, related to Windows NT, is service Pack 4. Additional technical specifications as well as the latest drivers and firmware levels can be found at:

<http://www.jni.com>

The JNI adapters listed in Table 1-1 on page 2 support the Point-to-Point, Arbitrated Loop, and Switched Fabric topologies.

Note: FC tape devices are currently not supported by all JNI Fibre Channel HBAs mentioned in the Table 1-1 on page 2.

1.3 Supported SAN fabric components

The SAN fabric components that are supported in the Windows NT and 2000 environment are covered in the following sections:

- ▶ IBM 2109 portfolio
- ▶ INRANGE
- ▶ McDATA

1.3.1 General information

There are two types of Gigabit Interface Converters (GBIC) and three types of fiber optic cable commonly used in Fibre Channel environments, as described in Table 1-2.

Table 1-2 Common Fibre Channel cabling configurations

GBIC type	Cable type	Distance (1 Gb/s)	Distance (2 Gb/s)
Longwave	9 micron, singlemode	10 km	10 km
Shortwave	50 micron, multimode	500 m	300 m
Shortwave	62.5 micron, multimode	300 m	150 m

Shortwave GBICs and 50 micron cable are usually used to connect devices to the fabric, since most device adapters only support shortwave GBICs, as well as for short inter-switch links (ISLs). Longwave GBICs and 9 micron cable are used when required for extended distance.

Attenuation

Fibre Channel is a high speed communication technology. Because of this, it is very sensitive to link attenuation. The attenuation is even more of a critical factor when connection speeds rise from 1 Gb/s to 2 Gb/s and even higher.

To limit the effects of attenuation, we recommend that you:

- ▶ Have as few connections as possible on a single link
- ▶ Make sure that all connectors are clean
- ▶ Make sure that all connectors are properly seated
- ▶ Make sure that you only use one type of fiber on a single link (never mix 50 micron and 62.5 micron fiber)

Zoning

We strongly recommend that you use zoning to:

- ▶ Enhance the security of your Fibre Channel fabric
- ▶ Avoid interactions between different hosts
- ▶ Secure the storage devices that do not support LUN masking, like tape drives

1.3.2 The IBM 2109 family

Current members of the IBM 2109 family of switches are described in Table 1-3.

Table 1-3 The IBM 2109 family

Model	Ports	Speed	Connectors	Rack space
IBM TotalStorage SAN Switch S08	8	1 Gb/s	SC	1U
IBM TotalStorage SAN Switch S16	16	1 Gb/s	SC	2U
IBM TotalStorage SAN Switch F16	16	1 or 2 Gb/s	LC	1U

Features and limitations

The IBM 2109 family of switches support both fabric-attached and loop-attached devices on any port, and switch domain IDs 1-239 are supported.

Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on requirements.

1.3.3 INRANGE

The INRANGE switches currently sold by IBM (the IBM 2042 family) are described in Table 1-4.

Table 1-4 The INRANGE switches

Model	Ports	Speed	Connectors	Rack space
2042-001 INRANGE FC/9000 Fibre Channel Director	64	1 Gb/s	SC	20U
2042-128 INRANGE FC/9000-128 Fibre Channel Director	128	1 Gb/s	SC	40U

In addition, IBM also sells the INRANGE FC/9000 40U cabinet (2042-C40). The cabinet is able to hold one 128-port switch or two 64-port switches.

Features and limitations

The INRANGE FC/9000 family of switches support both fabric-attached and loop-attached devices on any port.

Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on requirements.

1.3.4 McDATA

The McDATA switches currently sold by IBM (the IBM 2031 and 2032 families) are described in Table 1-5.

Table 1-5 The McDATA switches

Model	Ports	Speed	Connectors	Rack space
2031-016 McDATA ES-3016 Fabric Switch	16	1 Gb/s	LC	1U
2031-032 McDATA ES-3032 Fabric Switch	32	1 Gb/s	LC	1.5U
2031-L00 McDATA ES-1000 Loop Switch	8+1	1 Gb/s	SC	1U
2032-001 McDATA ED-5000 Enterprise Fibre Channel Director	32	1 Gb/s	SC	18U
2032-064 McDATA ED-6064 Enterprise Fibre Channel Director	64	1 Gb/s	LC	9U

In addition, IBM also sells the McDATA FC-512 Cabinet (2032-C36).

Features and limitations

The only McDATA switch that supports loop attached devices (tape drives, for example) is the McDATA ES-1000 Loop Switch.

The ports in the 2032-064 can be added in cards. Each card contains four ports. There are three different cards available:

- ▶ Four shortwave (multimode) ports
- ▶ Four longwave (singlemode) ports
- ▶ Three shortwave ports and one longwave port

On the 2031-016 and 2031-032 switches, any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC.

The 2031-L00 switch has two different types of ports:

- ▶ Eight switching hub ports that can only be configured with shortwave (multimode) GBIC (or copper interface).
- ▶ One bridge port, for connecting the loop switch to the fabric, that can be configured with a shortwave (multimode) or longwave (singlemode) GBIC.

The McDATA switches support fabrics of up to 31 switches.

1.4 Storage

The IBM Fibre Channel attached storage systems supported by the Windows 2000 and Windows NT operating systems are listed in Table 1-6 and Table 1-7.

Table 1-6 IBM Fibre Channel attached disk storage systems

Storage system	Operating systems
IBM TotalStorage FAStT200	Windows NT 4.0: <ul style="list-style-type: none"> ▶ Service Pack 6a or greater Windows 2000 <ul style="list-style-type: none"> ▶ Service Pack 1 or greater
IBM TotalStorage FAStT500	Windows NT 4.0 (Server and Enterprise edition): <ul style="list-style-type: none"> ▶ Service Pack 6a or greater Windows 2000 (Server and Advanced server) <ul style="list-style-type: none"> ▶ Service Pack 2 or greater
IBM TotalStorage FAStT700	Windows NT 4.0 (Server and Enterprise edition): <ul style="list-style-type: none"> ▶ Service Pack 6a or greater Windows 2000 (Server and Advanced server) <ul style="list-style-type: none"> ▶ Service Pack 2 or greater
IBM TotalStorage 2106 MSS	Windows NT 4.0 (Server and Enterprise edition) Windows 2000 (Server and Advanced server)
IBM TotalStorage 2105 ESS	Windows NT 4.0 (Server and Enterprise edition) ▶ Windows 2000 (Server and Advanced server)

Table 1-7 IBM Fibre Channel attached tape storage systems

Storage system	Notes
LTO 3584	Windows NT Server 4.0 <ul style="list-style-type: none"> ▶ service Pack 6a or later Windows 2000 <ul style="list-style-type: none"> ▶ Build 2195 or later
Magstar 3590 drives	Windows NT Server 4.0 <ul style="list-style-type: none"> ▶ service Pack 6a or later Windows 2000 Build 2195 or later

For additional information related to IBM TotalStorage disk and tape storage devices portfolio, contact your local IBM Storage Specialist.

Windows 2000 and Windows NT may also support storage systems produced by other vendors, but you should check the vendor's Web site to see if their storage equipment supports the Windows environment.

1.5 FC host bus adapter card installation procedure

In the following sections, we cover in a step-by-step manner the driver installation procedure for QLogic, Emulex and JNI FC host bus adapter cards. We use a Netfinity server 5500 running Windows NT SP6 and Windows 2000 Advanced Server SP2.

1.5.1 Preparing to install the Fibre Channel HBAs

There are a number of common steps that you should do before installation of Fibre Channel HBAs:

- ▶ Check if you have a free and supported PCI slot for the Fibre Channel HBAs.
- ▶ Download the latest driver.
- ▶ Download the latest adapter firmware.

The latest driver and firmware versions for all adapter types can be found at the following Web sites:

<http://www.qlogic.com>
<http://www.emulex.com>
<http://www.jni.com>

However, any special considerations for specific Fibre Channel HBAs are mentioned in installation sections.

1.5.2 Installing QLogic Fibre Channel HBAs family on Windows 2000

The QLogic SANblade 2200 and 2300 family of Fibre Channel HBAs (unlike the QLA 2100 model) support FC-AL as well as Point-To-Point and Switched Fabric connection. This section lists the steps you must follow in order to install the cards and the appropriate driver on Windows 2000 operating systems, as well as the configuration utility for specific cards.

Installation procedure

SANblade 2200 and 2300 adapters are plug-and-play devices that are automatically detected by Windows 2000. Perform the following steps in order to install QLA2200F or QLA2300F drivers:

1. Shut down the system.

2. Install the QLogic adapter card in the host system.
3. Turn on the system.
4. Windows 2000 detects QLA2200 or QLA2300 and starts **Found New Hardware Wizard**.
5. Select **Search for a suitable driver for my device** (recommended) and click **Next**.
6. Select **SCSI and RAID controllers** in the Hardware Types list, then click **Next**.
7. Click **Have Disk**. The **Install From Disk** dialog box appears.
8. Insert the disk containing the QLogic driver into an appropriate drive, then type the location of the driver.
9. Click **OK**.
10. Select the desired QLogic Fibre Channel HBAs, then click **Next**. Windows 2000 indicates that the wizard is ready to install the device.
11. Click **Next**.
12. If the **Digital Signature Not Found** dialog box appears, click **Yes**.
13. Click **Finish**.
14. Remove the disk that contains the QLogic driver from the drive, then click **Yes** to restart the system.

1.5.3 Installing QLogic Fibre Channel HBAs family on Windows NT

This section lists the steps you have to follow in order to install card and appropriate driver on Windows NT operating systems.

Installation procedure

1. Shut down the system.
2. Install the QLogic adapter card in the host system.
3. Turn on the system.
4. On the Windows NT taskbar, select **Start, Settings** and **Control Panel**.
5. Double-click the **SCSI Adapters** icon. The **SCSI Adapter** window displays.
6. Select the **Drivers** tab and then click **Add**.
7. Click **Have Disk**.
8. Enter the path to the Driver's location and click **OK**.
9. Select the driver and select **OK** to install necessary files.
10. Click **Yes** to restart the system.

1.5.4 QLogic configuration utility

QLogic provides a few utilities which can help in configuring and monitoring Fibre Channel adapter cards. We cover QLview which provides configuration information, statistics, diagnostics and alarm notification for adapter as well as for attached storage devices. You can download it from the following Web site:

<http://www.qlogic.com>

In this section we discuss the QLview stand-alone configuration used for monitoring QLogic adapters locally. The installation procedure is straightforward. QLview for Fibre Setup is a self-extracting program that installs the QLview for Fibre application and related software. Just double-click the **QLview Fibre Setup** icon and follow the wizard's instructions. More detail with respect to this can be found at the QLogic Web site.

After a successful installation, run the QLview utility and the following screen appears, as shown in Figure 1-4.

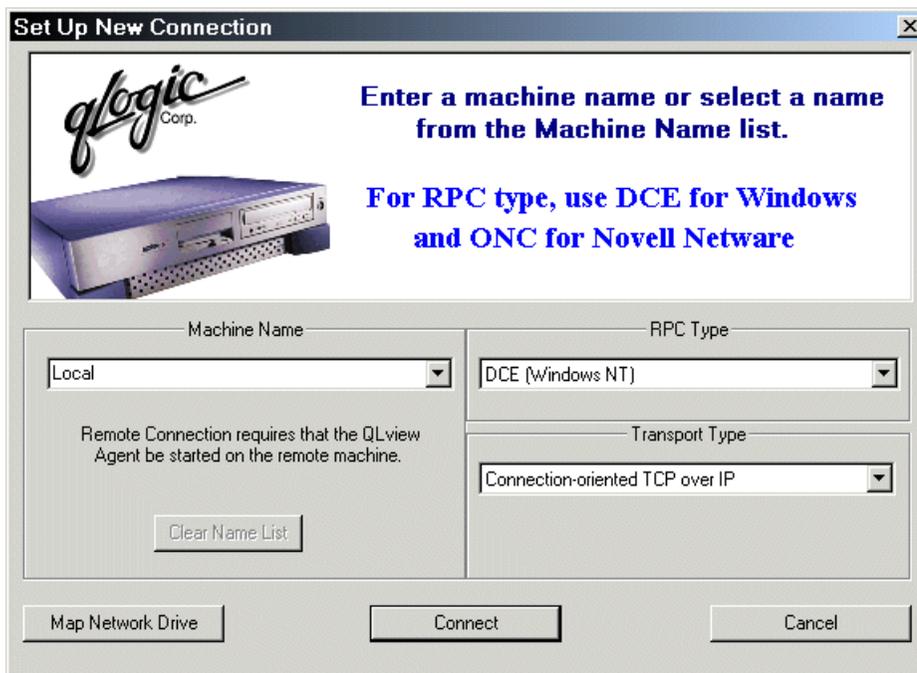


Figure 1-4 QLview for Fibre utility — Set Up Connection screen

The **Map Network Drive** button lists the name of all computers currently on the network to assist you in finding the names of the machines. In order to get more information related to QLogic Fibre Channel adapter card, click **Connect** as shown in Figure 1-5.

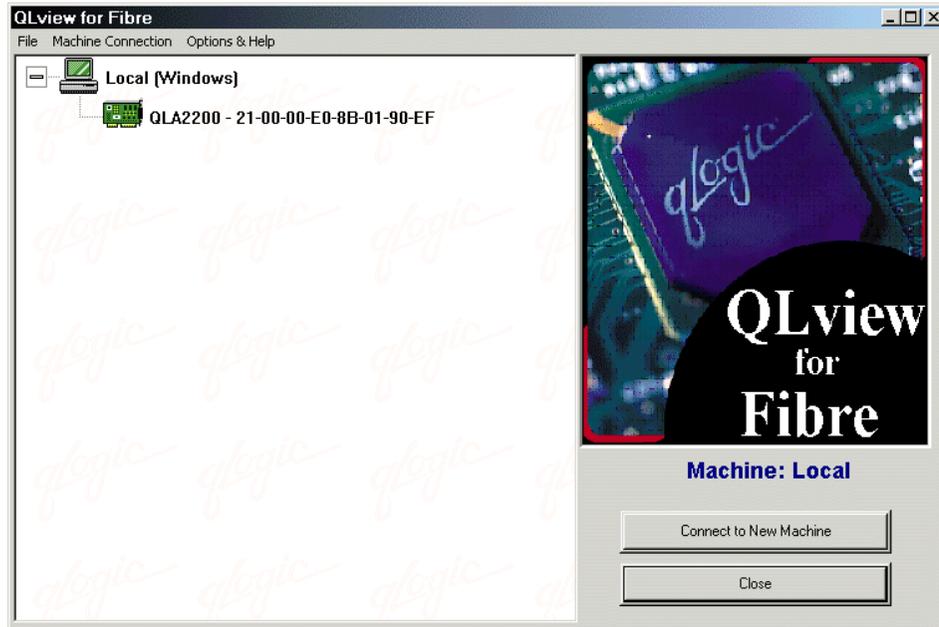


Figure 1-5 QLview for Fibre utility — Main Window

We can assume that the QLogic Fibre Channel HBA is successfully installed and connected to the SAN switch or direct to a storage device, unless you get the screen shown in Figure 1-6.

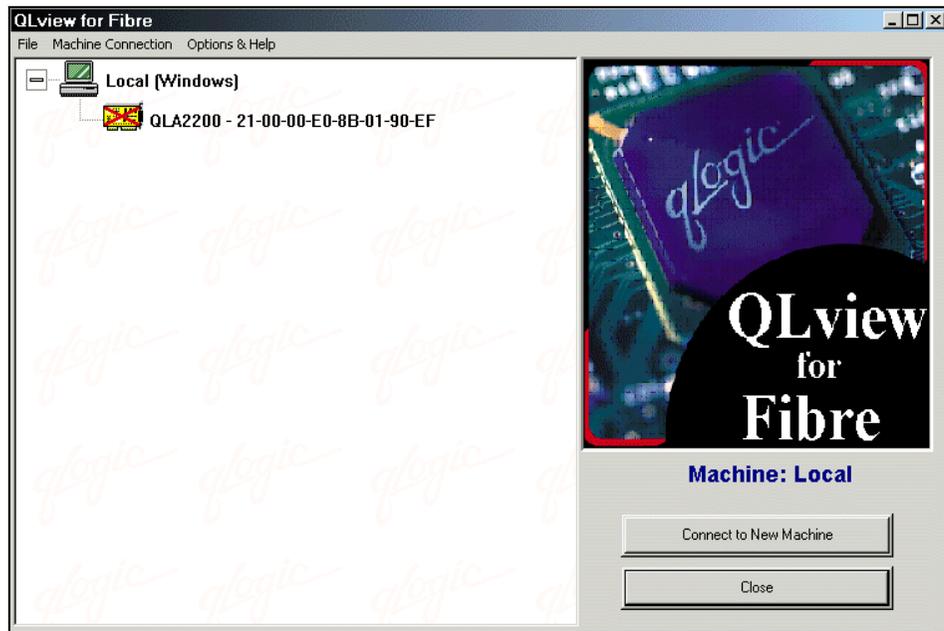


Figure 1-6 QLview for Fibre utility

In this case, the Fibre Channel HBAs is not connected to any storage device or SAN switch.

Select a specific QLogic adapter. On the right side of the window, you can find the adapter WWN as well as BIOS, firmware, and driver version. See Figure 1-7.

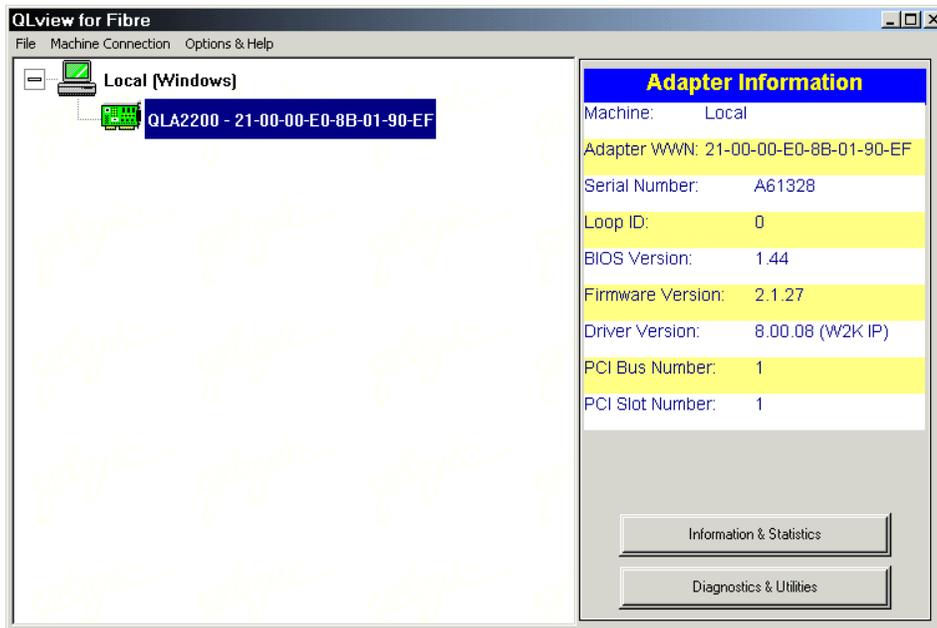


Figure 1-7 QLview for Fibre utility — adapter information

In addition to the QLview for Fibre Utility, there is also the *QLogic Fast Utility*, which is used for changing adapter configuration settings. For instance, if you want to boot from an FC disk drive attached to the adapter, use the *QLogic Fast Utility* in order to enable it. You can invoke it by pressing the **Alt** and **Q** keys (when requested) during Windows boot. We recommended that you carefully read the QLogic Fast Utility user guide before changing adapter settings.

1.5.5 Installing Emulex Fibre Channel HBAs in Windows 2000

This section lists the steps you have to follow in order to install the card and appropriate driver on Windows 2000 operating systems. Windows 2000 Service Pack 2 is recommended.

Installation procedure

1. Shut down the system.
2. Install the Emulex Fibre Channel HBAs in the host system.
3. Turn on the system.
4. A window displays **Found New Hardware Wizard**. Click **Next**.
5. Select **Display a list of known drivers for this device so that I can choose a specific driver**.

6. Click **Next**.
7. Click **Have Disk** and insert the containing driver to the appropriate drive then type the location of the driver.
8. Click **OK**.
9. Select the Emulex Fibre Channel card and then click **Next**. Windows 2000 indicates that the wizard is ready to install the device.
10. Click **Yes** to restart the system.

The adapter contains two green and one yellow LEDs. Immediately after successful installation of the adapter and the appropriate driver, the green LED in the middle is ON indicating that the adapter has power. As soon as the adapter is connected to a SAN switch or to a storage device directly, both green LED's are ON as well as the yellow one which signifies port activity. This is shown in Figure 1-8 on page 18.

1.5.6 Installing Emulex in Windows NT

The appropriate driver for both LP8000 and LP9000 adapters family are installed by default during your installation of Windows NT, as well as the Light Pulse utility (see "Emulex configuration utility" on page 18). You should update the existing driver to the latest version, which can be found at:

<http://www.emulex.com>

We use the 4.53a5 Emulex driver version. Windows NT Service Pack 3 is the minimum operating system requirement.

Installation procedure

1. Shut down the system.
2. Install the Emulex Fibre Channel HBAs in the host system.
3. Turn on the system.
4. On the Windows NT taskbar, select **Start, Settings** and **Control Panel**.
5. Double-click the **SCSI Adapters** icon. The **SCSI Adapter** window displays.
6. Select the **Drivers** tab and then click **Add**.
7. Click **Have Disk**.
8. Enter the path to the Driver's location and click **OK**.
9. Select the driver and select **OK** to install necessary files.
10. Click **Yes** to restart the system.

The adapter contains two green LEDs and one yellow LED. Immediately after successful installation of the adapter and the appropriate driver, the green LED in the middle is ON, indicating that the adapter has power. As soon as the adapter is connected to the SAN switch or to a storage device directly, both green LEDs are ON, as well as the yellow one which signifies port activity, as shown in Figure 1-8.

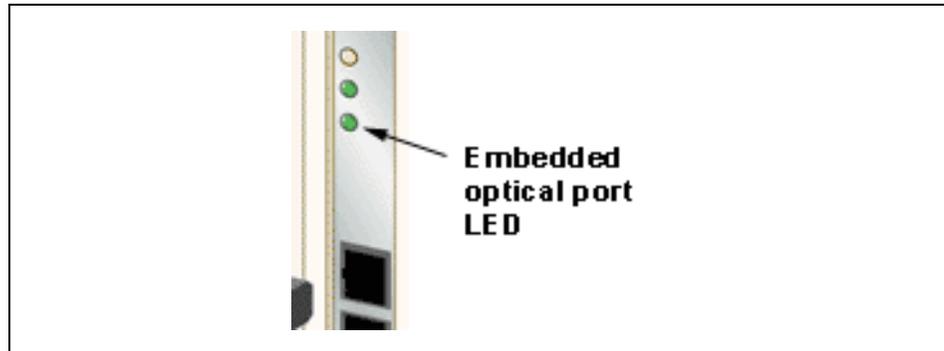


Figure 1-8 Emulex LP8000 Fibre Channel HBAs — Status LEDs

1.5.7 Emulex configuration utility

During the Emulex Fibre Channel HBAs driver installation, Light Pulse utility installation occurs. Light Pulse is a driver utility for Windows NT and Windows 2000 which allows you to modify the registry and upgrade adapter firmware. In addition to this, you can view adapter information and status.

In order to run the Light Pulse utility, click the **Start** button and click **Run**. In the Open section write down **Lputilnt** and the following window appears as shown in Figure 1-9.

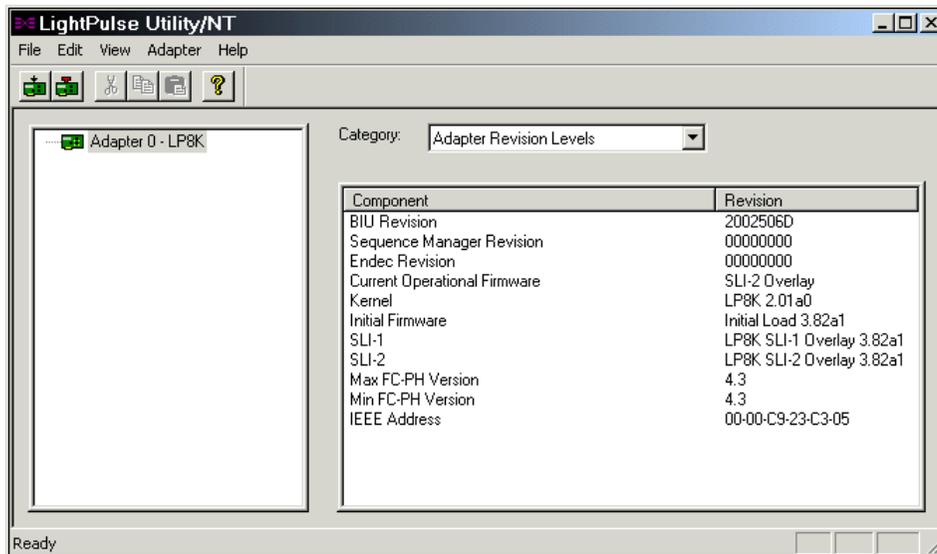


Figure 1-9 Light Pulse utility — main window

By selecting the appropriate adapter and changing the **Category** section, you can obtain much useful information about the driver itself, as well as performing other actions on it. Scroll up the **Category** section and select the desired action, as shown in Figure 1-10.

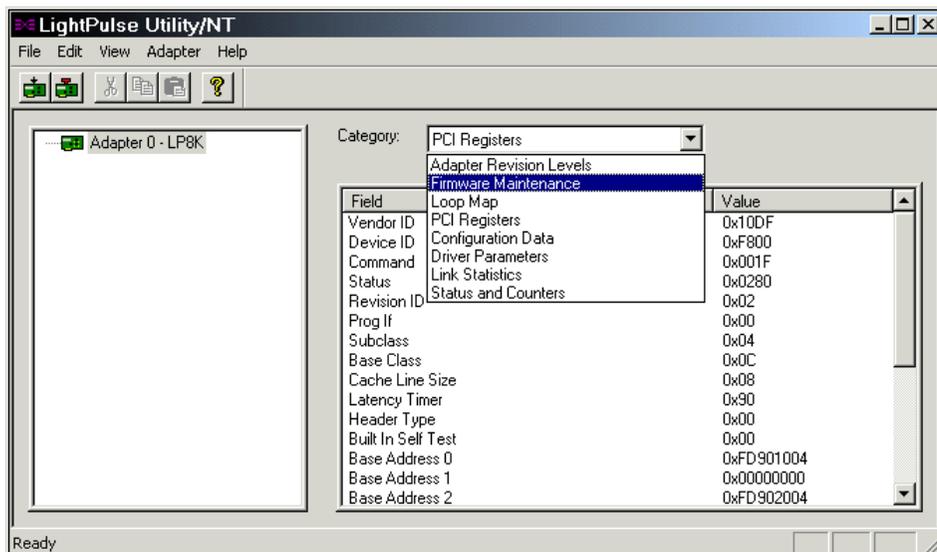


Figure 1-10 LightPulse Utility — Category list

It is recommended to use Light Pulse in order to download the latest version of adapter Firmware. Select **Firmware Maintenance** under the **Category** list and click the **Download** button, as shown in Figure 1-11.

Note: Do not upgrade the firmware if it is not suggested by the vendor. Check the vendor site for the latest news related to the firmware level. Read carefully all instructions before any action is taken.

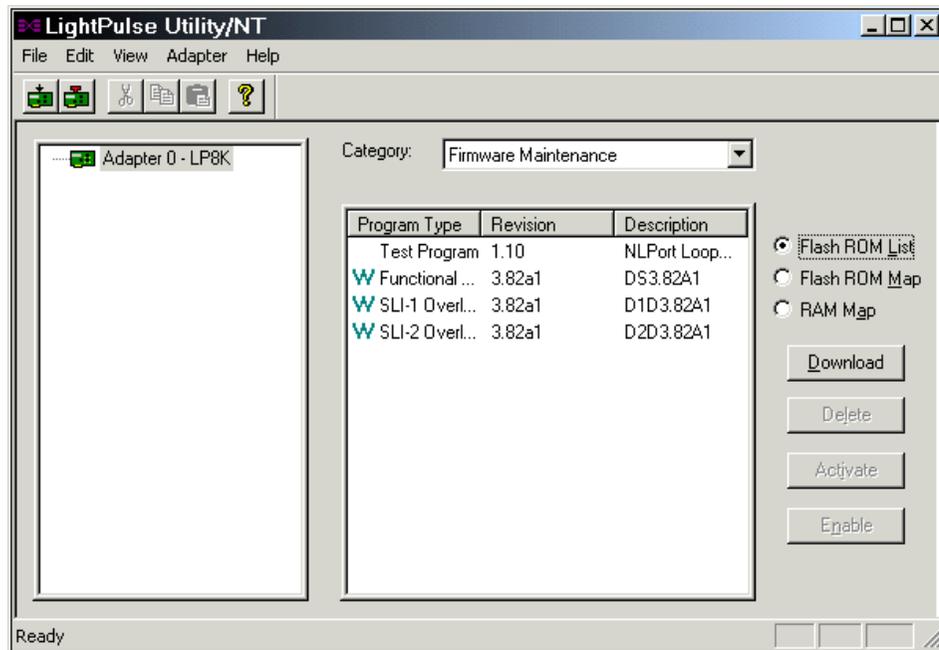


Figure 1-11 LightPulse Utility — Firmware download

At this point, enter the appropriate firmware destination path and click **Open**. The new software is transferred to the host bus adapter. In order to verify that the new firmware level is active, expand the **Category** list and select **Adapter Revision Levels**, as shown in Figure 1-9 on page 19.

1.5.8 Installing JNI Fibre Channel host bus adapters in Windows 2000

Microsoft Windows 2000 has a default built-in driver, the *JNI 1160 Fibre Channel Adapter*, which supports FCE-3210/6410 and FCE2-6410/6412 Fibre Channel HBAs. However, it is highly recommended to update your current installed driver to the latest version.

Installation procedure

1. Shut down the system.
2. Install JNI adapter card in the host system.
3. Turn on the system.
4. Windows 2000 automatically detects a new added Fibre Channel HBAs and installs the default JNI driver (already built-in Windows itself).
5. Right-click **My Computer** icon and select **Properties**.
6. Select **Hardware** tab and click **Device Manager**.
7. Double-click the **SCSI and RAID controllers** to expand the view.
8. Double-click **JNI 1160 Fibre Channel Adapter** and then click the **Driver** tab.
9. Click **Update Driver**. The **Upgrade Device Driver Wizard** appears.
10. Click **Next** to start the update operation.
11. Select **Display a list of the known drivers for this device** and then click **Next**.
12. Click **Have Disk**.
13. Enter the path to the Driver's location and click **OK**.
14. Select appropriate JNI FCE adapter and click **Next**. Windows 2000 indicates that the wizard is ready to install the device.
15. Click **Next**.
16. If the **Digital Signature Not Found** dialog box appears, click **Yes**.
17. Click **Finish**.
18. Click **Yes** to restart the system.

Note: To see storage targets under Windows 2000, you must install the EZ Fibre Utility in order to map the target storage device (see "JNI configuration utility" on page 22).

1.5.9 Installing JNI Fibre Channel host bus adapters in Windows NT

Windows NT 4.0 SERVICE PACK 4 is required when installing FCE-3210 and FCE-6410. For FCE2-6410 and FCE2-6412, JNI recommends that you install Service Pack 5 or greater to ensure proper operation of your system.

Installation procedure

1. Shut down the system.
2. Install the adapter card in the host system.
3. Turn on the system.
4. On the Windows NT taskbar, select **Start, Settings,** and **Control Panel.**
5. Double-click the **SCSI Adapter** icon. The **SCSI Adapter** window displays.
6. Select the **Drivers** tab and then click **Add.**
7. Click **Have Disk.**
8. Enter the path to the Driver's location and click **OK.**
9. Select the driver and select **OK** to install necessary files.
10. Click **Yes** to restart the system.

Note: To see storage targets under Windows NT, you must install the EZ Fibre Utility in order to map the target storage device (see 1.5.10, "JNI configuration utility" on page 22.)

1.5.10 JNI configuration utility

In order to configure and monitor adapter status, we recommend that you install the EZ Fibre Configuration Utility for Windows NT and Windows 2000. You can download it from:

<http://www.jni.com>

It is found under the drivers section. The minimum Windows NT requirement for EZ Fibre Configuration Utility implementation is Service Pack 6.

The EZ Fibre Configuration utility provides following features:

- ▶ Displays the JNI adapters installed in a system
- ▶ Displays the Fibre Channel targets and devices currently visible to each HBA
- ▶ Simplifies configuring JNI driver parameters
- ▶ Detects configuration errors
- ▶ Saves and restores adapter configuration information

In addition to the aforementioned functions, the EZ Fibre Configuration Utility is also used for mapping storage devices. When ready to connect Fibre Channel HBAs to a storage device directly or via fabrics, both Windows 2000 and Windows NT users should use the following mapping storage device procedure:

1. In the EZ Fibre Configuration Utility - Main Menu, right-click the adapter and select **Clear LLZ Targets**. Then click **Yes**.
2. With the adapter selected, click **LUN Level Zoning**.
3. Click the **Mapped** column and select the target LUN you want to be mapped.
4. Click **Commit Changes** at the bottom of the screen.
5. Restart the system.

1.6 Installing the storage system

In this section we describe a Windows based server connection to a storage device after successful Fibre Channel HBAs installation. We connected a Netfinity 5500 to an IBM TotalStorage 2106 Modular Storage Server via an IBM 2109 S16 SAN Switch. Procedures for storage definition and allocation, as well as switch zones creation, are not described here, since they differ among various manufacturers. Read carefully any documentation that comes with your storage system.

IBM 2109 S16 connectivity check

Before you try to allocate and define the storage device in Windows 2000 or Windows NT, you should check if the SAN switch is ready. Assume that the switch is already initialized and configured. That means that it is configured in the local network with an IP address. Find out the switch IP address and see if it is reachable from your Windows 2000 or Windows NT environment by typing the **ping <switch_ip_address>** command using the Command Prompt interface.

If you get a reply from the switch, you can access it using the telnet session as well as a Web browser. If you want to access the switch through a telnet session, as soon as the connection is established, you will be prompted to enter a login name and password.

Administrator privileges are necessary in order to make some changes in the switch configuration files. The default administrator login name is **admin**, and the password is **password**. At that point you should be familiar with switch interface commands in order to gather switch information or to create a zone. For more information about switch commands and their syntax, type **help** from the switch command line. To check if your FC HBA is properly connected to a switch port, type the **switchshow** command, as shown in Example 1-1.

Example 1-1 Switch connectivity check

Fabric OS (tm) Release a2.4.1c

login: **admin**

Password:

itsosw3:admin> **switchShow**

switchName: itsosw3

switchType: 2.4

switchState: Online

switchRole: Principal

switchDomain: 3

switchId: fffc03

switchWwn: 10:00:00:60:69:10:64:cf

switchBeacon: OFF

port 0: id Online F-Port 21:00:00:e0:8b:02:a7:2d

port 1: -- No_Module

port 2: id Online E-Port 10:00:00:60:69:20:1d:4e "itsosw1" (downstream)

port 3: -- No_Module

port 4: id No_Light

port 5: id Online F-Port 10:00:00:00:c9:23:c3:05

port 6: id Online E-Port 10:00:00:60:69:20:1d:4e "itsosw1"

port 7: id No_Light

port 8: sw Online F-Port 10:00:00:00:c9:22:bf:75

port 9: -- No_Module

port 10: sw No_Light

port 11: sw No_Light

port 12: sw Online L-Port 1 public

port 13: id No_Light

port 14: id Online F-Port 10:00:00:00:c9:25:8e:54

port 15: id No_Light

itsosw3:admin> **exit**

Connection closed by foreign host.

The online status of the specific port, fabric attachment (F_Port), and FC HBA WWN next to it, are the confirmation that the FC HBA has performed a fabric login to the switch.

An alternative way to access the switch is by using a Web browser, such as Internet Explorer. Just type the switch IP address and the switch main menu appears as shown in Figure 1-20.

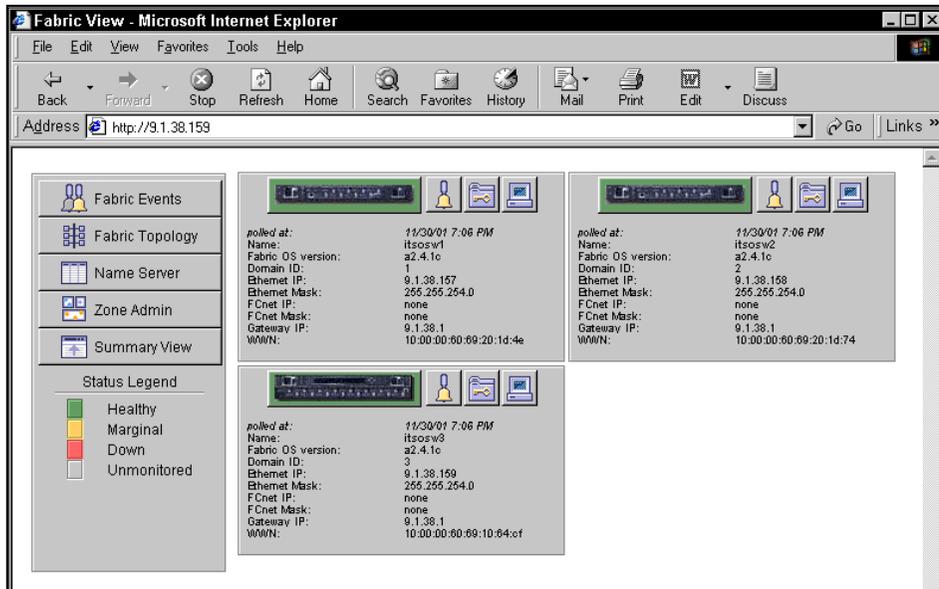


Figure 1-12 IBM 2109 Switch — configuration main menu

Click the switch to which your server is attached. The window shown in Figure 1-13 appears.



Figure 1-13 IBM 2109-S16 Switch Management Application

Click the switch port where the Windows server is connected, and the basic port information is displayed, as shown in Figure 1-14.

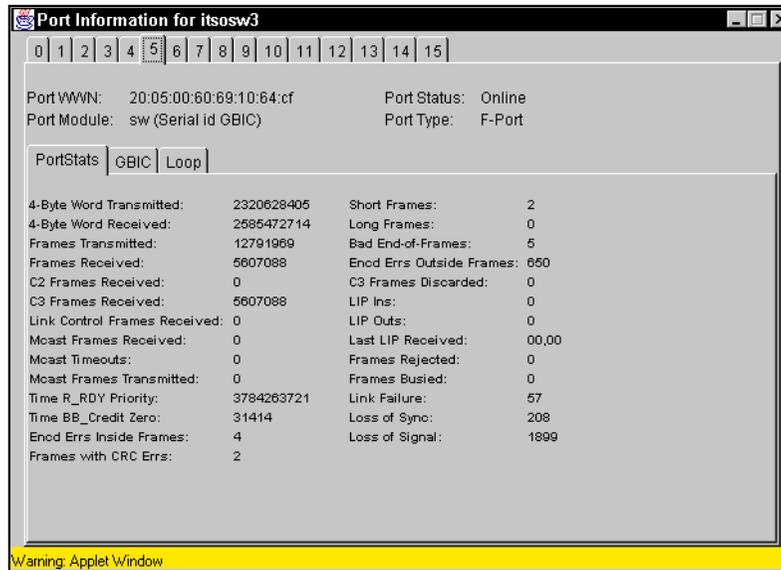


Figure 1-14 IBM 2109-S16 — Port Information

This is the easiest way to find out all the relevant switch information and to configure your own zone before you access the storage device from a Windows environment. Additionally, you do not have to know the switch login name and password in order to check if a switch is correctly connected to the server.

The switch connectivity check procedure is different for each switch vendor. For INRANGE switches, you should use the IN-VSN Enterprise Manager application to find the same information. For McDATA switches, you can find the information in the EFC Manager application.

Windows 2000 environment

As soon as you have connected your server to the storage device via a switch, defined and allocated a storage system, and created adequate switch zones, Windows recognizes the newly added storage device. You have to reboot your Windows 2000 or Windows NT in order to use the new storage device. Immediately after rebooting, follow these steps:

1. Select **Start, Program, Administrative Tool, and Computer Management**.
2. Under **Storage** tree, click **Disk Management** folder.
3. **Write Signature and Upgrade Disk Wizard** window appears, as shown in Figure 1-15.



Figure 1-15 Write Signature and Upgrade Disk Wizard

4. Click **Next**.

5. Select a newly added disk in order to write a signature, as shown in Figure 1-16.

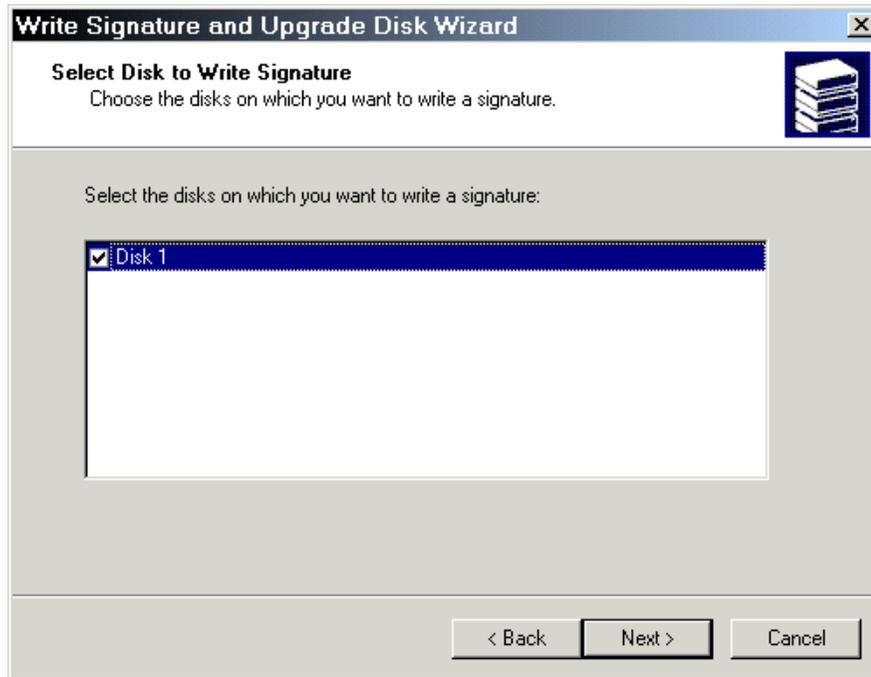


Figure 1-16 Select Disk

6. Click **Next**.

7. Click **Finish** to complete the Write Signature task, as shown in Figure 1-17.

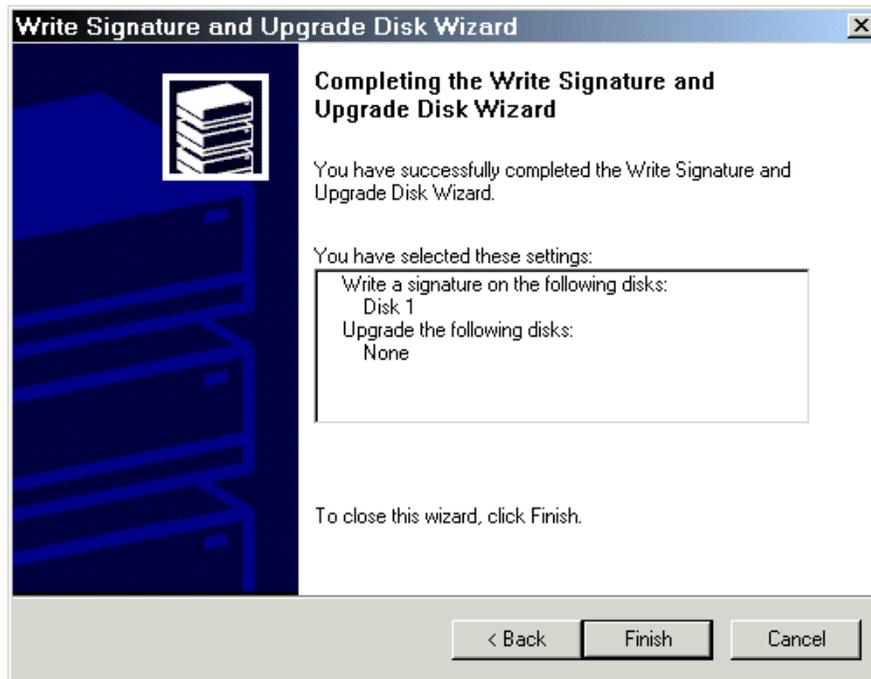


Figure 1-17 Completing the Write Signature task

- At that point, Windows 2000 recognizes the newly added storage devices, and it is able to perform a volume format, as shown in Figure 1-18.

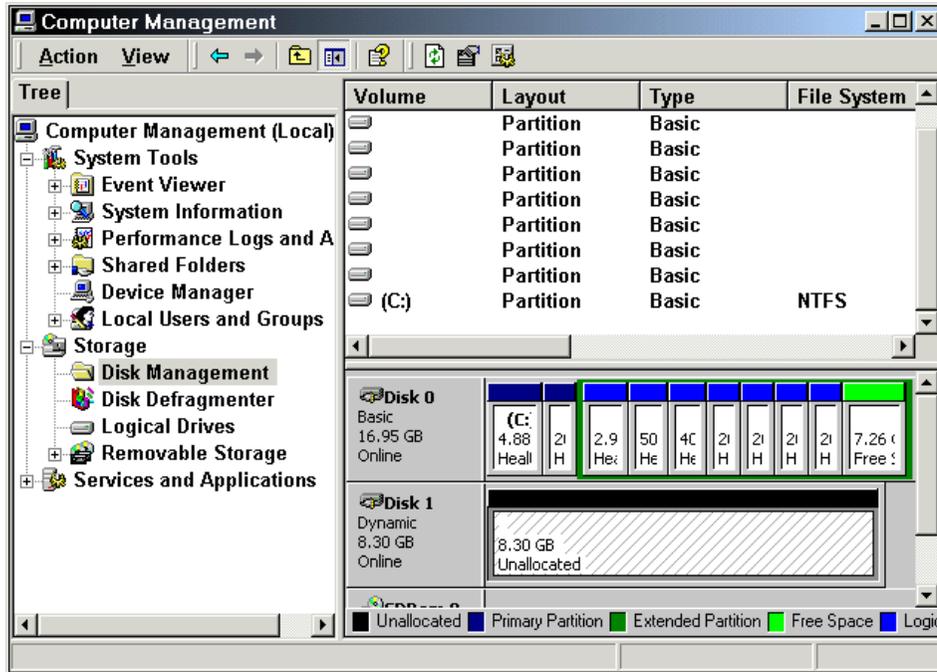


Figure 1-18 Disk Management — disk 1 unallocated and unformatted

Windows NT environment

These are the steps for the Windows NT environment:

- Select **Start, Program, Administrative Tool** and **Disk Administrator**, and the information appears, as shown in Figure 1-19.

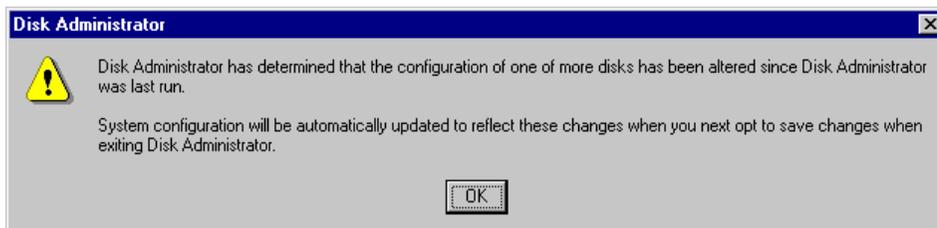


Figure 1-19 Disk Administrator change warning

- Click **OK**.

- Windows NT recognizes the newly added storage device, as shown in Figure 1-20.

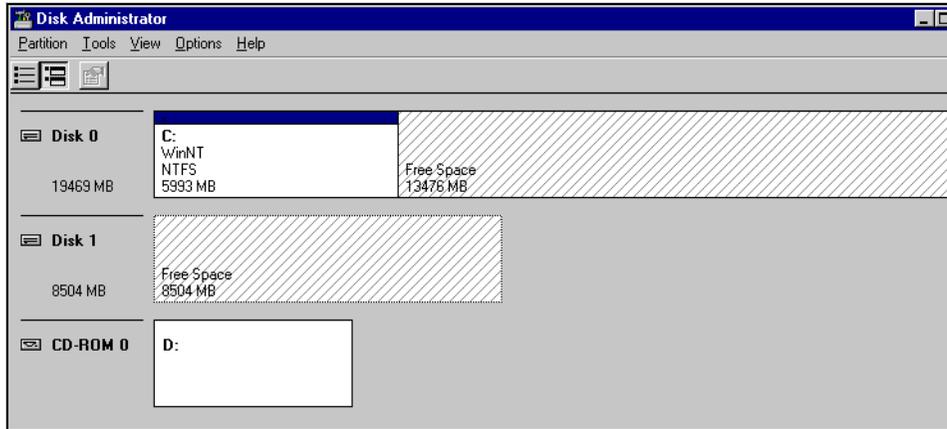


Figure 1-20 New device recognized

- At this point you can create one or more partitions on the newly added volume, as shown in Figure 1-21.

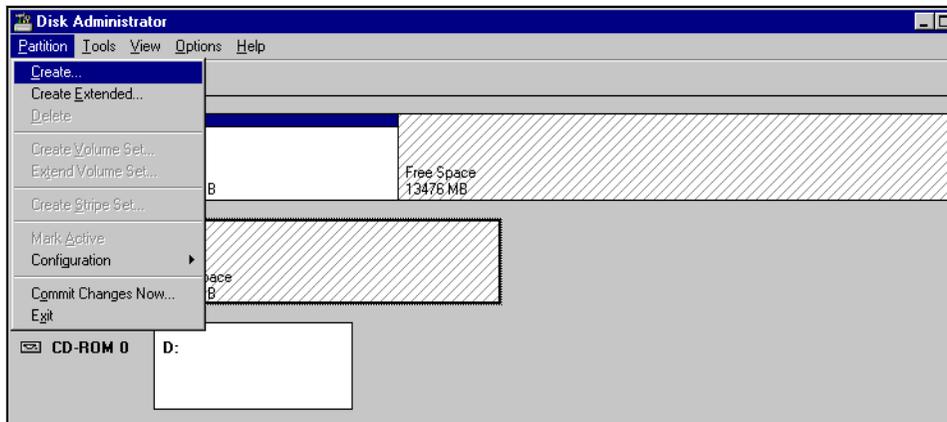


Figure 1-21 Creating partitions

5. Your newly added volume has a status, as shown in Figure 1-22.

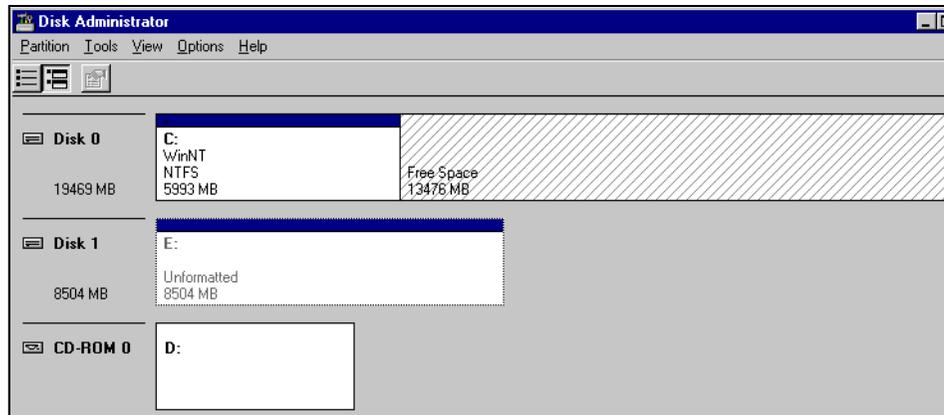


Figure 1-22 After creation of partitions

6. Prior to the format volume task, you should commit all changes which have been made, as shown in Figure 1-23.

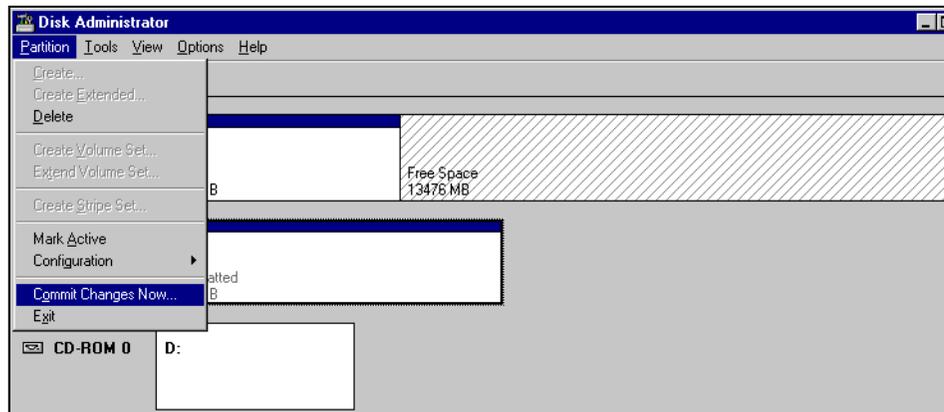


Figure 1-23 Commit changes

7. You may now use the newly added volume and create new partitions on it.



eServer pSeries Fibre Channel implementation

In this chapter we discuss how to successfully connect to a SAN fabric for the IBM eServer pSeries, RS/6000, and RS/6000 SP platforms. We cover the step-by-step procedure from start to finish, and the supported products which work with the eServer pSeries systems. We also point out where the latest firmware and device drivers can be obtained.

The following topics are discussed:

- ▶ “Introduction” on page 34
- ▶ “Fibre Channel host bus adapters” on page 34
- ▶ “SAN fabrics” on page 36
- ▶ “Storage” on page 40
- ▶ “Installation procedures” on page 40

2.1 Introduction

IBM AIX is the primary operating system used on the IBM eServer pSeries and RS/6000 systems. There are two versions of AIX currently available, AIX V4.3.3 and AIX 5L V5.1. Fibre Channel operations are generally the same on these operating systems, and we discuss any differences in their respective sections throughout this chapter.

The procedures shown in this chapter are also valid for attaching the nodes in an IBM RS/6000 SP system to a Fibre Channel network.

In this chapter, we only cover a single-path environment, where each fibre device can be accessed via only one path. If you have a multi-path environment, you need to use a storage system specific multi-path software, for example IBM Subsystem Device Driver for ESS.

Linux is also available for selected models of eServer pSeries and RS/6000 systems, but since there are no drivers for the Fibre Channel adapters available, we do not discuss it.

2.2 Fibre Channel host bus adapters

In this topic we will discuss the Fibre Channel host bus adapters (FC HBAs) applicable to the aforementioned products and operating systems.

2.2.1 IBM FC HBAs

There are two Fibre Channel host bus adapters available from IBM for the eServer pSeries and RS/6000 families, as specified in Table 2-1.

Table 2-1 IBM FC HBAs for IBM eServer pSeries and RS/6000

Adapter	f/c	Link speed	Connector
Gigabit Fibre Channel Adapter	6227	1 Gb/s	SC
2 Gigabit Fibre Channel Adapter	6228	1 or 2 Gb/s	LC

The device drivers for both of these adapters are included in the AIX install media. The 2 Gigabit adapter is the one we currently recommend. Both adapters have a single Fibre Channel port and support attaching both disk and tape devices.

Note: IBM RS/6000 model S70 does not support f/c 6228. On those systems, you must use f/c 6227 instead.

There are system-specific limitations on the number of host bus adapters you can install in a host. The limitations, as well as the recommended slots, are defined in *PCI Adapter Placement Reference*, SA38-0583, delivered with your system. The most recent version can also be found at:

http://www-1.ibm.com/servers/eserver/pseries/library/hardware_docs/index.html

You can download adapter firmware, AIX fixes and maintenance levels from:

<http://techsupport.services.ibm.com/server/fixes?view=pSeries>

You can also order the fixes on a CD at the same address.

2.2.2 Other FC HBAs

Other FC HBAs available for IBM eServer pSeries and RS/6000 servers are listed in Table 2-2.

Table 2-2 Other FC HBAs available for IBM eServer pSeries and RS/6000

Adapter	No of ports	Link Speed	Connector
Emulex LP6000	1	1 Gb/s	SC
Emulex LP7000E	1	1 Gb/s	SC
Emulex LP8000	1	1 Gb/s	SC
Emulex LP8000DC	2	1 Gb/s	SC
Emulex LP9000	1	1 Gb/s	SC
Emulex LP9002L	1	1 or 2 Gb/s	LC
Emulex LP9002DC	2	1 or 2 Gb/s	LC
Emulex LP9402DC	2	1 or 2 Gb/s	LC
JNI FCE-6410	1	1 Gb/s	SC
JNI FCE-3210	1	1 Gb/s	SC
JNI FCE2-6410	2	1 Gb/s	SC
JNI FCE2-6412	2	1 Gb/s	SC
JNI FCE-6460	1	1 or 2 Gb/s	LC

These adapters are supported by their manufacturer. For more information about the adapters and their limitations, see the manufacturer's Web site:

<http://www.emulex.com/>

<http://www.jni.com/>

2.3 SAN fabrics

The IBM eServer pSeries and RS/6000 hosts support two Fibre Channel network topologies:

- ▶ Fibre Channel Arbitrated Loop (FC-AL)
- ▶ Switched fabric

The IBM eServer pSeries and RS/6000 systems support no more than one host in an FC-AL loop. This means that FC-AL is only useful for connecting the host directly to a storage device, such as the IBM Enterprise Storage Server (ESS). FC-AL also has severe performance and scalability limitations. Because of these limitations, we only cover switched fabric topology. However, we also discuss attaching storage devices supporting only FC-AL attachment to the fabrics.

We cover products from the following three vendors:

- ▶ IBM 2109
- ▶ INRANGE
- ▶ McDATA

2.3.1 General information

There are two types of Gigabit Interface Converters (GBIC) and three types of fiber optic cable commonly used in Fibre Channel environments, as described in Table 2-3.

Table 2-3 Common Fibre Channel cabling configurations

GBIC type	Cable type	Distance (1 Gb/s)	Distance (2 Gb/s)
longwave	9 micron, singlemode	10 km	10 km
shortwave	50 micron, multimode	500 m	300 m
shortwave	62.5 micron, multimode	300 m	150 m

Shortwave GBICs and 50 micron cable are usually used to connect devices to the fabric, since most device adapters only support shortwave GBICs, as well as for short inter-switch links. Longwave GBICs and 9 micron cable are used when required for extended distance.

Attenuation

Fibre Channel is a high speed communication technology. Because of this, it is very sensitive to link attenuation. The attenuation is even more of a critical factor when connection speeds rise from 1 Gb/s to 2 Gb/s and even higher.

To limit the effects of attenuation, we recommend making sure that:

- ▶ You have as few connections as possible on a single link.
- ▶ All connectors are clean.
- ▶ All connectors are properly seated.
- ▶ Only one type of fiber is used on a single link (never mix 50 micron and 62.5 micron fiber).

Zoning

We strongly recommend that you use zoning to:

- ▶ Enhance the security of your Fibre Channel fabric.
- ▶ Avoid interactions between different hosts.
- ▶ Secure storage devices that do not support LUN masking, like tape drives.

Since the implementation of zoning and even the terminology used varies between switch manufacturers, we do not discuss it in any detail.

2.3.2 The IBM 2109 family

The current members of the IBM 2109 family of switches are described in Table 2-4.

Table 2-4 The IBM 2109 family

Model	Ports	Speed	Connectors	Rack space
IBM TotalStorage SAN Switch S08	8	1 Gb/s	SC	1U
IBM TotalStorage SAN Switch S16	16	1 Gb/s	SC	2U
IBM TotalStorage SAN Switch F16	16	1 or 2 Gb/s	LC	1U

Features and limitations

The IBM 2109 family of switches support both fabric-attached and loop-attached devices on any port, and switch domain ids 1-239 are supported.

Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on customer requirements.

For more information, see *IBM SAN Survival Guide Featuring the IBM 2109*, SG24-6127.

2.3.3 INRANGE

The INRANGE switches currently supported by IBM (the IBM 2042 family) are described in Table 2-5.

Table 2-5 The INRANGE switches

Model	Ports	Speed	Connectors	Rack space
2042-001 INRANGE FC/9000 Fibre Channel Director	64	1 Gb/s	SC	20U
2042-128 INRANGE FC/9000-128 Fibre Channel Director	128	1 Gb/s	SC	40U

In addition, IBM also sells the INRANGE FC/9000 40U cabinet (2042-C40). The cabinet is able to hold one 128-port switch or two 64-port switches.

Features and limitations

The INRANGE FC/9000 family of switches support both fabric-attached and loop-attached devices on any port.

Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on customer requirements.

For more information, see *IBM SAN Survival Guide Featuring the INRANGE Portfolio*, SG24-6150.

2.3.4 McDATA

The McDATA switches currently supported by IBM (the IBM 2031 and 2032 families) are described in Table 2-6.

Table 2-6 The McDATA switches

Model	Ports	Speed	Connectors	Rack space
2031-016 McDATA ES-3016 Fabric Switch	16	1 Gb/s	LC	1U
2031-032 McDATA ES-3032 Fabric Switch	32	1 Gb/s	LC	1.5U
2031-L00 McDATA ES-1000 Loop Switch	8+1	1 Gb/s	SC	1U
2032-001 McDATA ED-5000 Enterprise Fibre Channel Director	32	1 Gb/s	SC	18U
2032-064 McDATA ED-6064 Enterprise Fibre Channel Director	64	1 Gb/s	LC	9U

In addition, IBM also sells the McDATA FC-512 Cabinet (2032-C36).

Features and limitations

The only McDATA switch that supports loop attached devices (tape drives, for example) is the McDATA ES-1000 Loop Switch.

The ports in the 2032-001 and 2032-064 can be added in cards. The cards are not compatible between the models. Each card contains four ports. There are three different cards available:

- ▶ Four shortwave (multimode) ports
- ▶ Four longwave (singlemode) ports
- ▶ Three shortwave ports and one longwave port

On the 2031-016 and 2031-032 switches, any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC.

The 2031-L00 switch has two different types of ports:

- ▶ Eight switching Hub ports that can only be configured with shortwave (multimode) GBIC (or copper interface)
- ▶ One bridge port, for connecting the loop switch to the fabric, that can be configured with a shortwave (multimode) or longwave (singlemode) GBIC.

The McDATA switches support fabrics of up to 31 switches only.

For more information, see *IBM SAN Survival Guide Featuring the McDATA Portfolio*, SG24-6149.

2.4 Storage

The IBM Fibre Channel attached storage systems supported on eServer pSeries and RS/6000 machines are described in Table 2-7 and Table 2-8.

Table 2-7 IBM Fibre Channel attached disk storage systems

Storage system	Notes
IBM TotalStorage Enterprise Storage Server	
IBM TotalStorage FASiT500 Storage Server	
IBM TotalStorage SAN Controller 160 with IBM 7133 disk subsystems	Only supports IBM 2109 switch fabrics

Table 2-8 IBM Fibre Channel attached tape storage systems

Storage system	Notes
IBM TotalStorage Enterprise Tape System 3590	Requires feature 3510 or 9510 on the drive
IBM 3584 UltraScalable Tape Library	Requires feature 1456 or 1466 on the drive Only supports IBM 2109 switch fabrics
IBM 3583 Ultrium Scalable Tape Library with SAN Data Gateway Module Feature 8005	
Tape drives attached via IBM 2108-G07 IBM Storage Area Network Data Gateway or IBM 2108-R03 IBM Storage Area Network Data Gateway Router	

Other storage system vendors may also support their storage systems on eServer pSeries and RS/6000 systems.

2.5 Installation procedures

In this topic we discuss the installation procedures.

2.5.1 Preparations to install the FC HBAs

These are the preparations necessary in order to install the FC HBAs:

- ▶ Check that you have a free, supported PCI slot for the FC HBAs.
- ▶ Check that you have the correct version and latest maintenance level of AIX installed:
 - At the time of writing, the latest maintenance level is 4.3.3.0.9
- ▶ Obtain the latest drivers:
 - For the Gigabit Fibre Channel Adapter (f/c 6227), you need the packages `devices.pci.df1000f7.*`
 - For the 2 Gigabit Fibre Channel Adapter (f/c 6228), you need the packages `devices.pci.df1000f9.*`
- ▶ Obtain the latest adapter firmware.

Download adapter firmware, AIX fixes and maintenance levels from:

<http://techsupport.services.ibm.com/server/fixes?view=pSeries>

The page opened by this URL is shown in Figure 2-1.

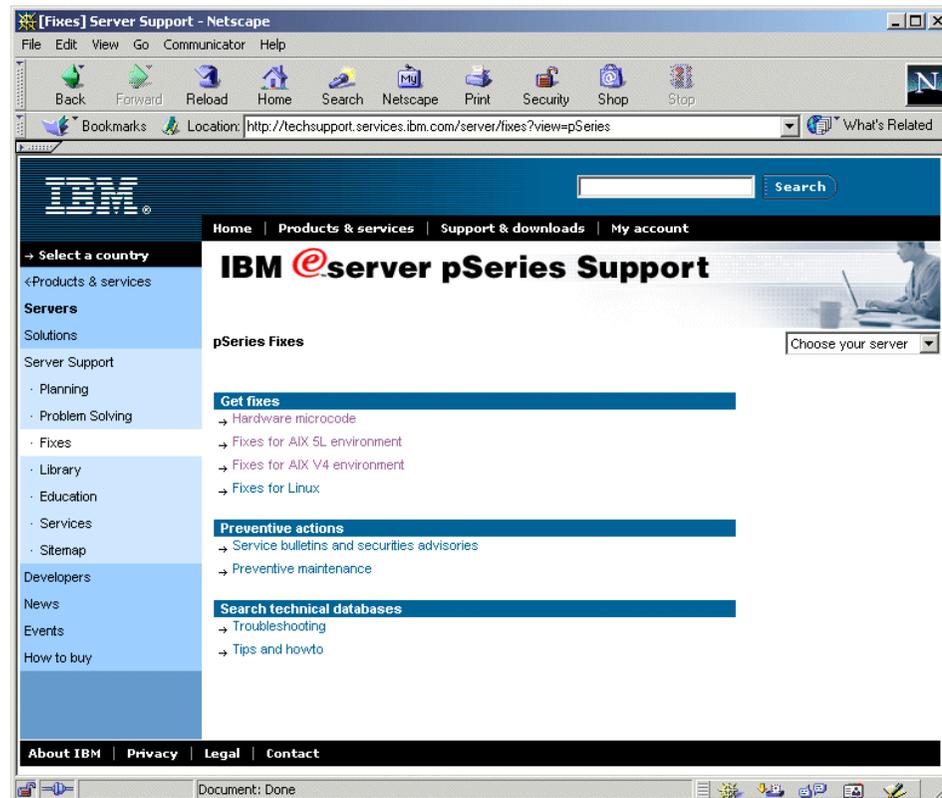


Figure 2-1 IBM eServer pSeries Support

We use the following hardware and software in our installation:

- ▶ IBM RS/6000 model F80, with AIX 4.3.3.0.9 installed
- ▶ IBM Gigabit Fibre Channel Adapter, f/c 6227
- ▶ AIX 4.3.3 maintenance level 9 packages in directory /install/aix43309
- ▶ Latest device driver updates for the Fibre Channel adapter, also located in directory /install/aix43309:
 - devices.pci.df1000f7.diag.4.3.3.76
 - devices.pci.df1000f7.com.4.3.3.76
 - bos.rte.lvm.4.3.3.76
- ▶ Latest microcode for the Fibre Channel adapter, version 02903290 (3.22A0)

2.5.2 Installing the FC HBAs

The following instructions explain how to install the FC HBAs:

1. Power on the system, if it is not already on, and log in as user root.
2. Check that you have the correct AIX maintenance level installed, as shown in Example 2-1.

Example 2-1 Checking the AIX maintenance level

```
# instfix -i|grep 4330-09_AIX_ML
  All filesets for 4330-09_AIX_ML were found.
#
```

3. Install the base level of device drivers for your FC HBAs.
 - a. Insert the install media into the CD-ROM drive, if you are installing from CD-ROM.
 - b. Start SMIT by typing:

```
# smitty devinst
```

You get the screen shown in Figure 2-2.

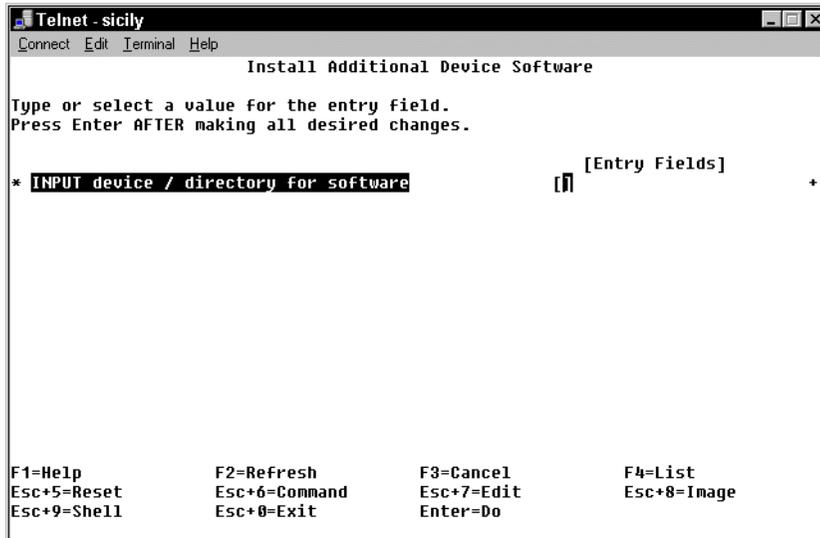


Figure 2-2 Install additional device software — select install medium

- c. Type in your install device or directory (`/dev/cd0`, for example) and press Enter

The next screen is shown in Figure 2-3.

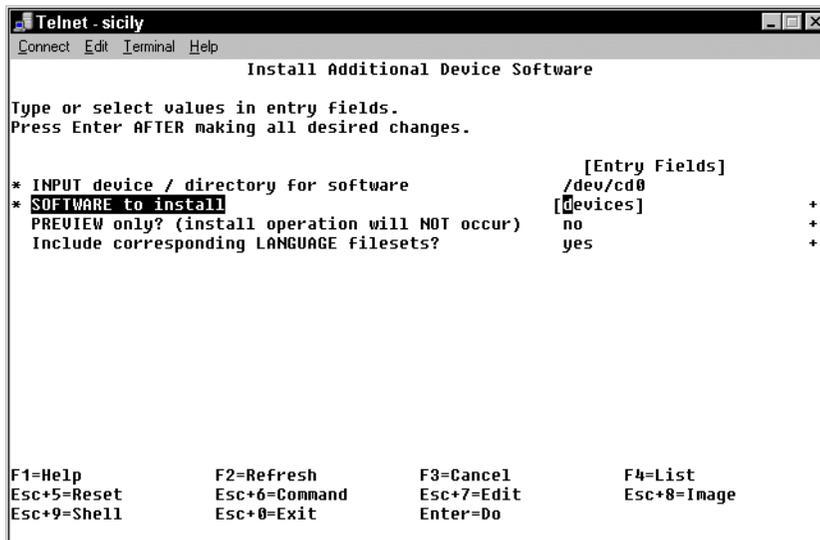


Figure 2-3 Install additional device software — initial settings

- d. Press F4 to get a list of the software packages.

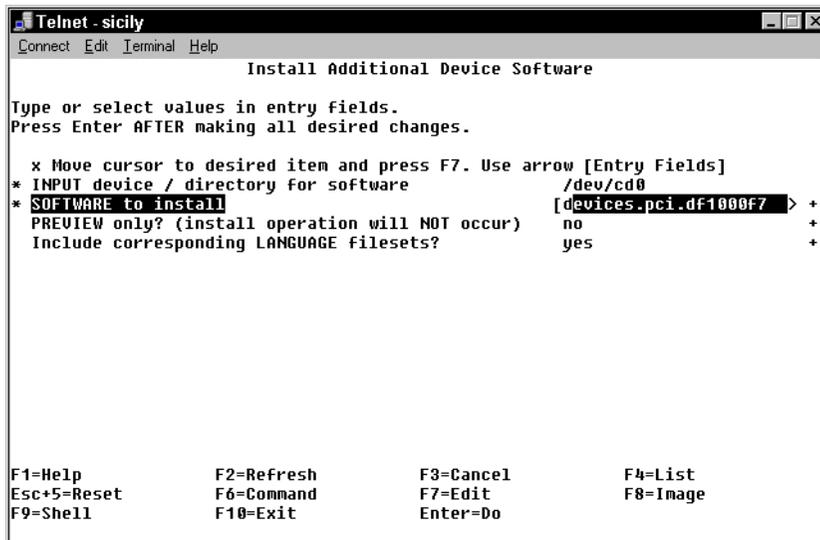


Figure 2-5 Install additional device software — installation settings

- h. Accept the choices by pressing Enter.
- i. You get the ARE YOU SURE window; press Enter to accept it.
- j. You get the COMMAND STATUS window, as shown in Figure 2-6.

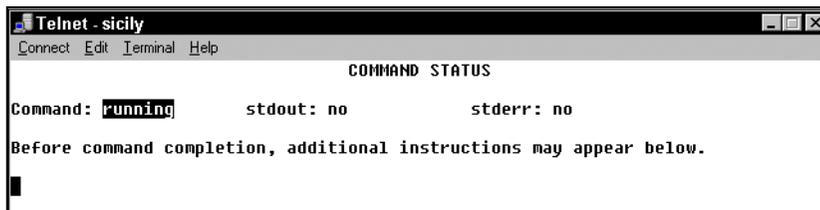


Figure 2-6 Install additional device software — installation in progress

- k. The system may ask you to mount another CD-ROM of the installation package, like that shown in Figure 2-7.

```

Telnet - sicily
Connect Edit Terminal Help

COMMAND STATUS

Command: running      stdout: yes      stderr: no

Before command completion, additional instructions may appear below.

[MORE...89]
US Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corp.

. . . . << End of copyright notice for devices.common.IBM.fc >>. . . .

Filesets processed: 4 of 6 (Total time: 13 secs).

installp: APPLYING software for:
          printers.msg.en_US.rte 4.3.3.0

installp: Please mount volume 3 on /dev/cd0.
          Press the enter key to continue.
          If volume 3 is not available you can enter "q" to terminate
          this process.

█
[BOTTOM]

```

Figure 2-7 Install additional device software — mount additional volume

- i. When the status changes from RUNNING to OK, scroll to the bottom of the screen, as shown in Figure 2-8, to check that installation ended in SUCCESS.

```

Telnet - sicily
Connect Edit Terminal Help

COMMAND STATUS

Command: OK          stdout: yes      stderr: no

Before command completion, additional instructions may appear below.

[MORE...154]
devices.pci.df1000f7.com 4.3.3.0      USR      APPLY      SUCCESS
devices.pci.df1000f7.rte 4.3.3.0      ROOT     APPLY      SUCCESS
devices.pci.df1000f7.com 4.3.3.0      ROOT     APPLY      SUCCESS
devices.common.IBM.fc.rte 4.3.3.0     USR      APPLY      SUCCESS
printers.msg.en_US.rte 4.3.3.0      USR      APPLY      SUCCESS
X11.msg.en_US.usm.rte 4.3.3.0     USR      APPLY      SUCCESS

installp: * * * A T T E N T I O N ! ! !
          Software changes processed during this session require this system
          and any of its diskless/dataless clients to be rebooted in order
          for the changes to be made effective.

█
[BOTTOM]

F1=Help          F2=Refresh      F3=Cancel1      F6=Command
F8=Image         F9=Shell        F10=Exit        /=Find
n=Find Next

```

Figure 2-8 Install additional device software — results

- m. Remove the installation media, if you are installing from CD-ROM.
Press F10 to exit SMIT.
4. Update the drivers you just installed to the latest versions.
 - a. Start SMIT by typing:

```
# smitty update_all
```

You will see the screen shown in Figure 2-9.

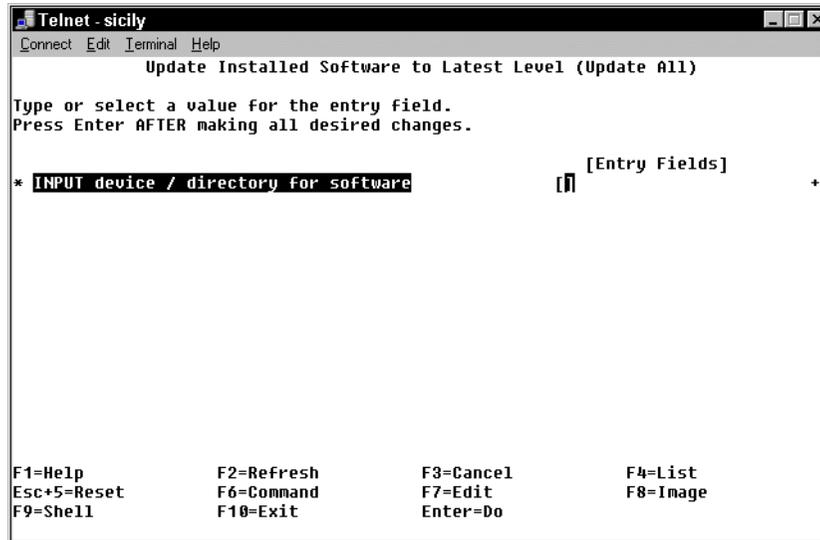


Figure 2-9 Update device software — install medium

- b. Type in your update directory or device (`/install/aix43309`, for example) and press Enter.

You see a screen like the one in Figure 2-10.

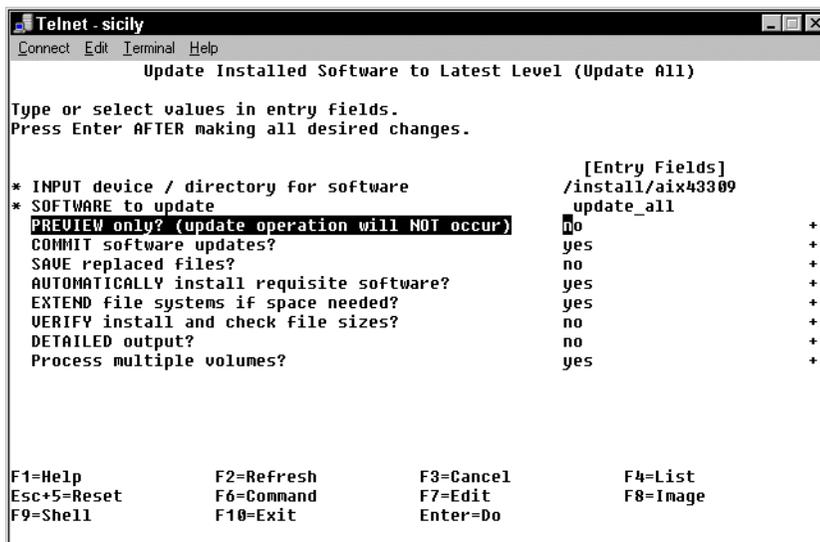


Figure 2-10 Update device software — installation settings

- c. Accept the screen by pressing Enter.
- d. You get the ARE YOU SURE window, press Enter to accept it.
Next you get the COMMAND STATUS window, as shown in Figure 2-11.

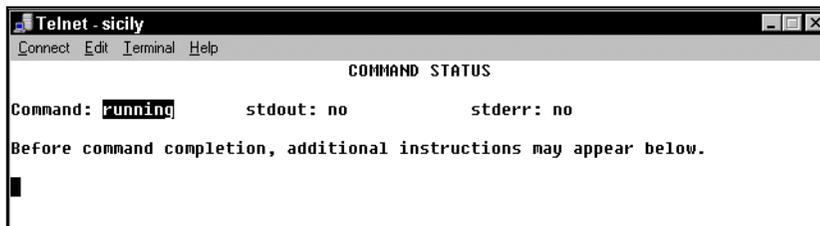


Figure 2-11 Update device software — installation in progress

- e. When the status changes from RUNNING to OK, scroll to the bottom of the screen, as shown in Figure 2-12, to check that installation ended in SUCCESS.

```

Telnet - sicily
Connect Edit Terminal Help

COMMAND STATUS

Command: OK          stdout: yes          stderr: no

Before command completion, additional instructions may appear below.

[MORE...179]
devices.pci.df1000f7.rte 4.3.3.75      USR      APPLY    SUCCESS
devices.common.IBM.fc.rte 4.3.3.75      USR      APPLY    SUCCESS
devices.common.IBM.fc.rte 4.3.3.75      USR      COMMIT   SUCCESS
devices.pci.df1000f7.com 4.3.3.76      USR      COMMIT   SUCCESS
devices.pci.df1000f7.rte 4.3.3.75      USR      COMMIT   SUCCESS
devices.pci.df1000f7.diag 4.3.3.76      USR      COMMIT   SUCCESS

installp: * * * A T T E N T I O N ! ! !
          Software changes processed during this session require this system
          and any of its diskless/dataless clients to be rebooted in order
          for the changes to be made effective.

[BOTTOM]

F1=Help          F2=Refresh      F3=Cancel      F6=Command
F8=Image         F9=Shell        F10=Exit       /=Find
n=Find Next

```

Figure 2-12 Update device software — results

- f. Press F10 to exit SMIT.
5. Shut down the system.
6. Install the host bus adapter card in a supported slot.

Depending on the model of your system, the adapter is installed by IBM's Service Support Representative (SSR), or you may need to install it yourself. You can find the installation instructions in the documentation provided with your system.

Note: Some eServer pSeries and RS/6000 models support hot-pluggable adapters. If your system is one of those models, you have the option of installing the adapter without shutting down the system. If you want to make use of the hot pluggable feature, follow the directions in the installation manual delivered with your system.

7. Connect the adapter to your Fibre Channel fabric.

Note: The Fibre Channel device drivers choose the network topology, when the adapter is configured. If your adapter is not connected to a Fibre Channel fabric, it will be configured as FC-AL attached. You can correct the problem by unconfiguring and reconfiguring the adapter as follows:

```
# rmdev -R1 fcs0
fscsi0 Defined
fcs0 Defined
# cfgmgr
#
```

8. Power on the system and log in as user root.
9. Check that the adapter is correctly configured as shown in Example 2-2.
You get a line for each of the Fibre Channel cards installed in the system.

Example 2-2 Checking the Fibre Channel adapter

```
# lsdev -C|grep fcs
fcs0      Available 3A-08      FC Adapter
#
```

10. Check the firmware level of your adapter, as shown in Example 2-3.

Example 2-3 Checking the firmware level

```
# lscfg -v1 fcs0

fcs0          3A-08          FC Adapter

Part Number.....09P1162
EC Level.....D
Serial Number.....KT05110105
Manufacturer.....0010
FRU Number.....09P1173
Network Address.....10000000C9258E54
ROS Level and ID.....02903290
Device Specific.(Z0).....4002206D
Device Specific.(Z1).....10020193
Device Specific.(Z2).....3001506D
Device Specific.(Z3).....02000909
Device Specific.(Z4).....FF101450
Device Specific.(Z5).....02903290
Device Specific.(Z6).....06113290
Device Specific.(Z7).....07113290
Device Specific.(Z8).....20000000C9258E54
Device Specific.(Z9).....SS3.22A0
Device Specific.(ZA).....S1F3.22A0
```

```
Device Specific.(ZB).....S2F3.22A0
Device Specific.(YL).....P1-I9/Q1
```

```
#
```

You can find the firmware level as the value of the "ROS Level and ID" parameter.

You can also find the World Wide Name (WWN) of the adapter as the value of the "Network Address" parameter. You need to know the WWN to configure your storage system.

11. If the firmware level is lower than the latest one, upgrade it according to the instructions provided with the firmware package
12. Check that the Fibre Channel adapter is using the correct protocol (al or switch) as shown in Example 2-4

Example 2-4 Checking the protocol

```
# lsattr -El fscsi0
scsi_id      0x31e00 Adapter SCSI ID          False
attach       switch How this adapter is CONNECTED False
sw_fc_class  3      FC Class for Fabric              True
#
```

You can find the protocol used as the value of the "attach" parameter.

You can also see whereabouts in the switch fabric our card is attached to by looking at the hexadecimal "scsi_id" parameter. From the rightmost byte, the first two numbers are always "00" for a switch-attached adapter, the third hexadecimal number (e) tells the port of the switch board, the fourth hexadecimal number (1) the number of the switch board within the switch, and the remaining one or two hexadecimal numbers (3) the domain id of the switch.

13. Optionally, check to see that the port is correctly configured on the switch
- The procedure you need to use is different for each switch vendor. We show the procedure for the IBM 2109 family of switches in Example 2-5. For INRANGE switches, you should use the IN-VSN Enterprise Manager application to find the same information. For McDATA switches, you can find the information in the EFC Manager application.

Example 2-5 Checking the FC HBA connectivity to the switch

```
# telnet 9.1.38.159
Trying...
Connected to 9.1.38.159.
Escape character is '^]'.

Fabric OS (tm) Release a2.4.1c

login: admin
Password:

itsosw3:admin> switchShow
switchName:    itsosw3
switchType:    2.4
switchState:   Online
switchRole:    Principal
switchDomain:   3
switchId:      fffc03
switchWwn:     10:00:00:60:69:10:64:cf
switchBeacon:  OFF
port 0: id Online      F-Port 21:00:00:e0:8b:02:a7:2d
port 1: -- No_Module
port 2: id Online      E-Port 10:00:00:60:69:20:1d:4e "itsosw1" (downstream)
port 3: -- No_Module
port 4: id No_Light
port 5: id No_Light
port 6: id Online      E-Port 10:00:00:60:69:20:1d:4e "itsosw1"
port 7: id No_Light
port 8: sw Online      F-Port 10:00:00:00:c9:22:bf:75
port 9: -- No_Module
port 10: sw No_Light
port 11: sw No_Light
port 12: sw Online     L-Port 1 public
port 13: id No_Light
port 14: id Online     F-Port 10:00:00:00:c9:25:8e:54
port 15: id No_Light
itsosw3:admin> quit
Connection closed.
#
```

From Example 2-5, we can see that the port we are using (port 14) is online. It is an F_Port (for fabric attachment), and the WWN listed is the same as that we see on our adapter as shown in Example 2-3 on page 50.

14. Your adapter is now correctly installed.

2.5.3 Installing the storage system

After the adapter installation, we have to create any required switch zones and allocate the storage for our host from the storage system we are using. As both of those steps are different for different switch manufacturers and different storage subsystems, they are not described here.

Note: The procedure described below is very generic. You may have to install additional device drivers depending on your storage system, like the Host Attachment Scripts for ESS or the Atape driver for tapes. Refer to the documentation that came with your storage system.

We have access to a single volume in a storage subsystem. We want to configure that volume to our host system. The procedure is as follows:

1. Try to run `cfgmgr` to scan for devices:

```
# cfgmgr  
#
```

2. If you get an error message like the one below, you have to install additional device drivers:

```
# cfgmgr  
cfgmgr: 0514-621 WARNING: The following device packages are required for  
device support but are not currently installed.  
devices.fcp.disk  
#
```

The procedure to install the new drivers is as follows:

- a. Start SMIT by typing

```
# smitty cfgmgr
```

You get the screen shown in Figure 2-13.

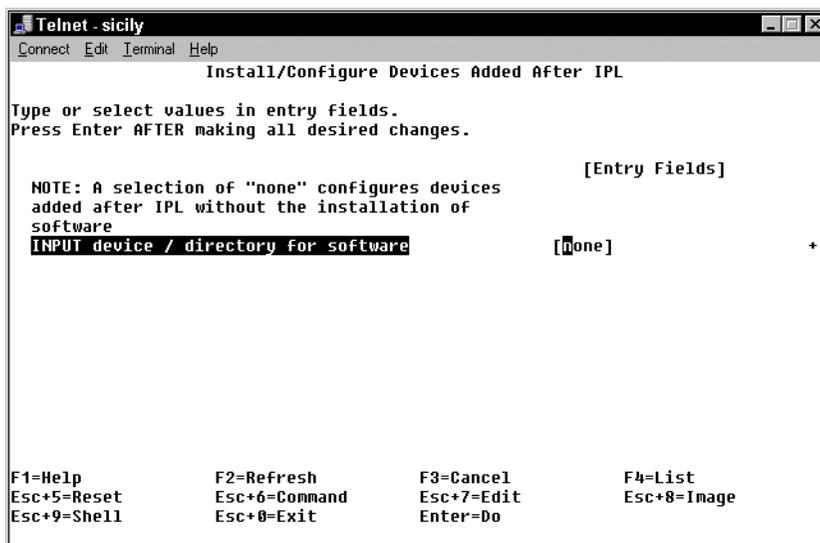


Figure 2-13 Install storage drivers — select install medium

- b. Choose or type in your install device or directory (`/dev/cd0`, for example).
- c. Accept the choices by pressing Enter.
- d. You get the COMMAND STATUS window as shown in Figure 2-14.

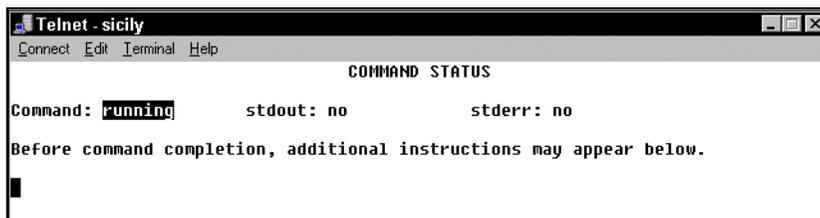


Figure 2-14 Install storage drivers — installation in progress

- e. When the status changes from RUNNING to OK, scroll to the bottom of the screen, as shown in Figure 2-15, to check that installation ended successfully.

```
Telnet - sicily
Connect Edit Terminal Help

COMMAND STATUS

Command: OK          stdout: yes          stderr: no

Before command completion, additional instructions may appear below.

[MORE...62]
. . . . . << End of copyright notice for devices.fcp.disk >>. . . .

Finished processing all filesets. (Total time: 11 secs).

install: * * * A T T E N T I O N ! ! !
        Software changes processed during this session require this system
        and any of its diskless/dataless clients to be rebooted in order
        for the changes to be made effective.

bosboot: Boot image is 8485 512 byte blocks.
█
[BOTTOM]

F1=Help          F2=Refresh      F3=Cancel      Esc+6=Command
Esc+8=Image      Esc+9=Shell     Esc+0=Exit     /=Find
n=Find Next
```

Figure 2-15 Install storage drivers — results

- f. Remove the installation media, if you are installing from CD-ROM.
- g. Press F10 to exit SMIT.
- h. To update the drivers you just installed to the latest versions, start SMIT by typing:

```
# smitty update_all
```

You see the screen shown in Figure 2-16.

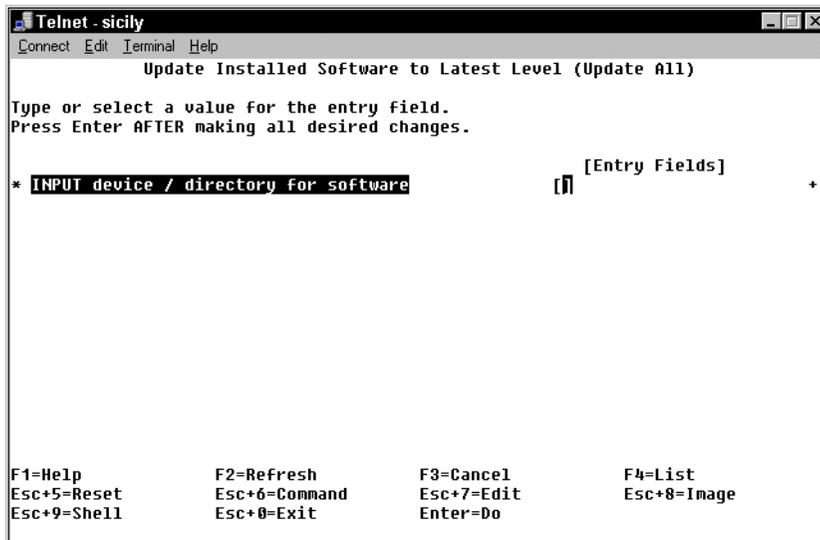


Figure 2-16 Update storage drivers — select install medium

- i. Type in your update directory or device (/install/aix43309, for example) and press Enter

You see a screen like the one in Figure 2-17.

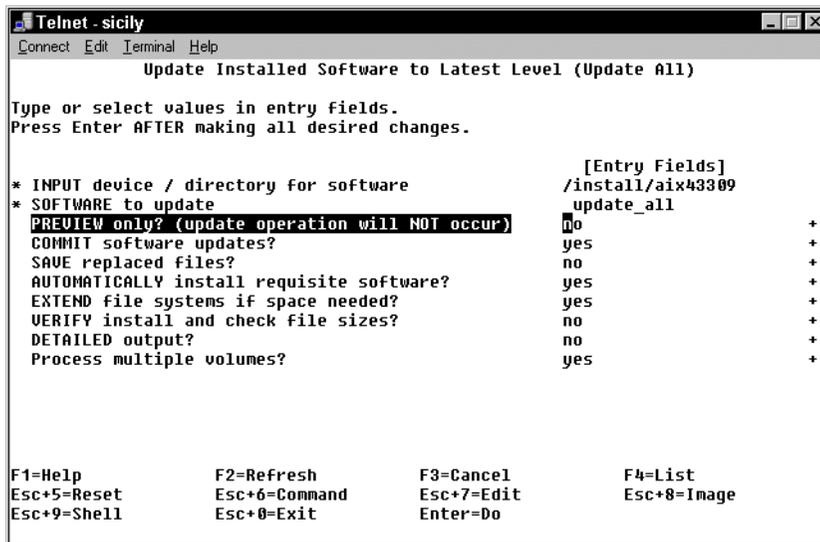


Figure 2-17 Update storage drivers — installation settings

- j. Accept the screen by pressing Enter.

k. You get the ARE YOU SURE window; press Enter to accept it.

You get the COMMAND STATUS window as shown in Figure 2-18.

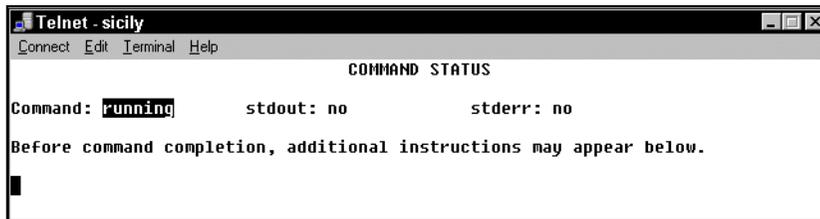


Figure 2-18 Update storage drivers — installation in progress

l. When the status changes from RUNNING to OK, scroll to the bottom of the screen, as shown in Figure 2-19, to check that installation ended in SUCCESS.

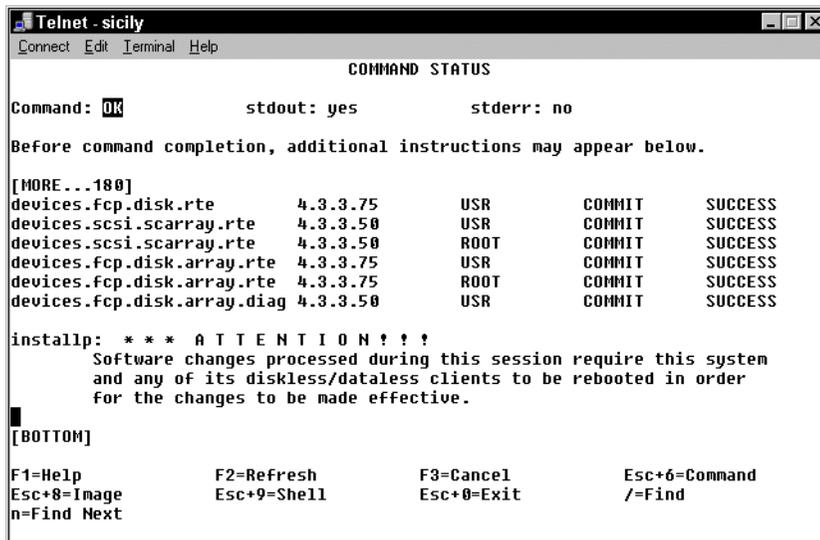


Figure 2-19 Update storage software — results

m. Press F10 to exit SMIT.

3. Check that the new storage device (disk in our case) is found:

```
# lsdev -Ccdisk
hdisk0 Available 11-08-00-2,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 11-08-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 3A-08-01 Other FC SCSI Disk Drive
#
```

4. Start using the device like any disk.

We use the procedure shown in Example 2-6 to create a volume group and filesystem on the new disk and test it.

Example 2-6 Creating and testing a filesystem

```
# mkvg -s 32 -y myfcvg hdisk2
0516-1254 mkvg: Changing the PVID in the ODM.
myfcvg
# crfs -v jfs -g myfcvg -m /fctest -A yes -a size=1048576
Based on the parameters chosen, the new /fctest JFS file system
is limited to a maximum size of 134217728 (512 byte blocks)

New File System size is 1048576
# mount /fctest
# cp /etc/hosts /fctest/hosts
# diff /etc/hosts /fctest/hosts
#
```

The connection is successful.



Sun Fibre Channel implementation

In this chapter we discuss how to successfully set up a SAN fabric for Sun Microsystems running Sun Solaris V2.6, 7 and 8. We cover the step-by-step procedure from start to finish, and the supported products which will work with these Sun systems.

The following topics are discussed:

- ▶ “Overview of the Sun system” on page 60
- ▶ “Sun supported Fibre Channel host adapters” on page 62
- ▶ “Sun supported Fibre Channel switches and hubs” on page 69
- ▶ “Supported storage on the Sun fabric” on page 74
- ▶ “Installation planning” on page 75
- ▶ “Installation procedure” on page 76
- ▶ “Device functionality checks” on page 90
- ▶ “Specific limitations” on page 91

3.1 Overview of the Sun system

There are a great number of Sun platform models which support the use of Fibre Channel to connect their system storage to them. For information on which ones support Fibre Channel, and in what configuration, we recommend you contact your Sun representative, or you can visit their Web site at the following URL:

<http://www.sun.com/products-n-solutions>

The Sun system can also have many different adapters, switches, and devices which it supports in certain defined configurations. These vary from product to product, and vendor to vendor. We strongly recommend that you contact your Sun representative or visit the varying vendor's Web pages for details on support. Usage of their products with the Sun platforms are documented on a "by-product" basis. Sun platforms support the standard PCI bus, or the Serial Bus (Sbus) for Fibre Channel host adapter (FC HBA) support. Sun supports Fibre Channel in the Arbitrated Loop (FC-AL), and Fibre Channel Switch (switched fabric) topologies.

With the FC-AL topology, Sun has a maximum limit of 50 LUNs per FC HBA. The following Sun platforms are supported by the broadest number of vendors for Fibre Channel support:

- ▶ Sun Sparc5, and Ultra1 Servers: with QLogic QL2100 adapter only.
- ▶ Sun Ultra2, 5, 10, 30, 60, 80, 220R, 250,420R and 450: with Sbus JNI FC64-1063, and JNI-1963 PCI, adapters. Additionally, there is a JNI 1460 PCI adapter currently under test for PCI use with these platforms.
- ▶ Sun E3000-6500 Servers: with PCI adapters.
- ▶ Sun E3000-10000 Servers: with Sbus adapters.

The following is a list of Sun operating systems which support Fibre Channel attachment:

- ▶ Solaris 2.6 (also known as SunOS 5.6)
- ▶ Solaris 7 (also known as SunOS 5.7)
- ▶ Solaris 8 (also known as SunOS 5.8)

Sun supports fabric environments that are both zoned and cascaded. For details on these capabilities, contact your Sun representative. Sun supports the IBM Subsystem Device Driver (SDD) multipathing driver product for the IBM TotalStorage Enterprise Storage Server (ESS), or you may use Veritas DMP. The IBM SDD comes in two separate varieties for Sun:

- ▶ Sun host systems which are 32-bit mode and/or running Solaris 2.6
- ▶ Sun host systems which are 64-bit mode, and running Solaris 7 or 8

With some other fibre devices, Sun platforms use Sun's Alternate Pathing (AP) to provide failover paths to its storage devices.

Sun has exercised a great deal of partnership to develop a large variety of vendor choices for customers in FC HBAs, switches, hubs, and storage devices which are supported in their SAN environments. We cover the ones tested and proofed in IBM environments in detail.

Required kernel patches for Fibre Channel support

In the Table 3-1 we show the minimum required kernel patches for running Fibre Channel with Solaris 2.6, 7, and 8. Ensure that you have these patches installed prior to trying to run a Fibre Channel environment.

Table 3-1 Solaris 2.6, 7, and 8 minimum patch levels for Fibre Channel

Solaris 2.6	Solaris 7	Solaris 8
105181-23 kernel update	106541-12 kernel update	108528-03 kernel update
105356-16 sd, ssd drivers	106924-06 isp driver	109524-02 ssd driver
105580-16 glm driver	106925-04 glm driver	109657-01 isp driver
105600-19 isp driver	107147-08 pci driver	108974-03 sd, uata drivers
Not applicable	107458-10 dad, sd, ssd, uata drivers	Not applicable

Note: Care should be taken when you this command. Frequently, patches are obsoleted by newer versions of the kernel update, which means you will no longer see the original patch number listed as an active patch at all. Such is the case with 105600-19 for the ISP driver, as it was rolled up into version 105181-29 of the kernel update on this system.

3.2 Sun supported Fibre Channel host adapters

Sun supports a great number of different vendor FC HBAs for their products, in both the 1 Gb/s, and 2 Gb/s models. An important point to consider here, when selecting the adapter you will use, is whether your other device vendors support your FC HBA.

A large selection of adapters are supported by some vendors, with only a few being documented as supported by all. We have listed the vendors, and their respective Web site URLs:

- ▶ QLogic Corporation
http://www.qlogic.com/products/fc_san_hostadapers.asp
- ▶ Emulex Corporation
<http://www.emulex.com/products/fc/index.html>
- ▶ JNI Corporation
<http://www.jni.com/Products/hba.cfm>
- ▶ Antares Microsystems
http://www.antares.com/fibre_channel/fibre_channel.htm
- ▶ Genroco, Inc.
<http://www.genroco.com/products/fcnet/index.html>
- ▶ ATTO Technology, Inc.
<http://www.attotech.com/fcha.html>

At the time of writing, IBM supports the following JNI, QLogic, and Emulex FC HBAs with which IBM has tested our products with the Sun platforms:

- ▶ Emulex LP8000 - PCI adapter: This requires the latest firmware and driver from Emulex. In Figure 3-1 we show this adapter. It supports “Open Boot” software on Solaris. All downloads, including documentation, are available at the following URL:
<http://www.emulex.com/ts/fc/docs/frame8kd.htm>

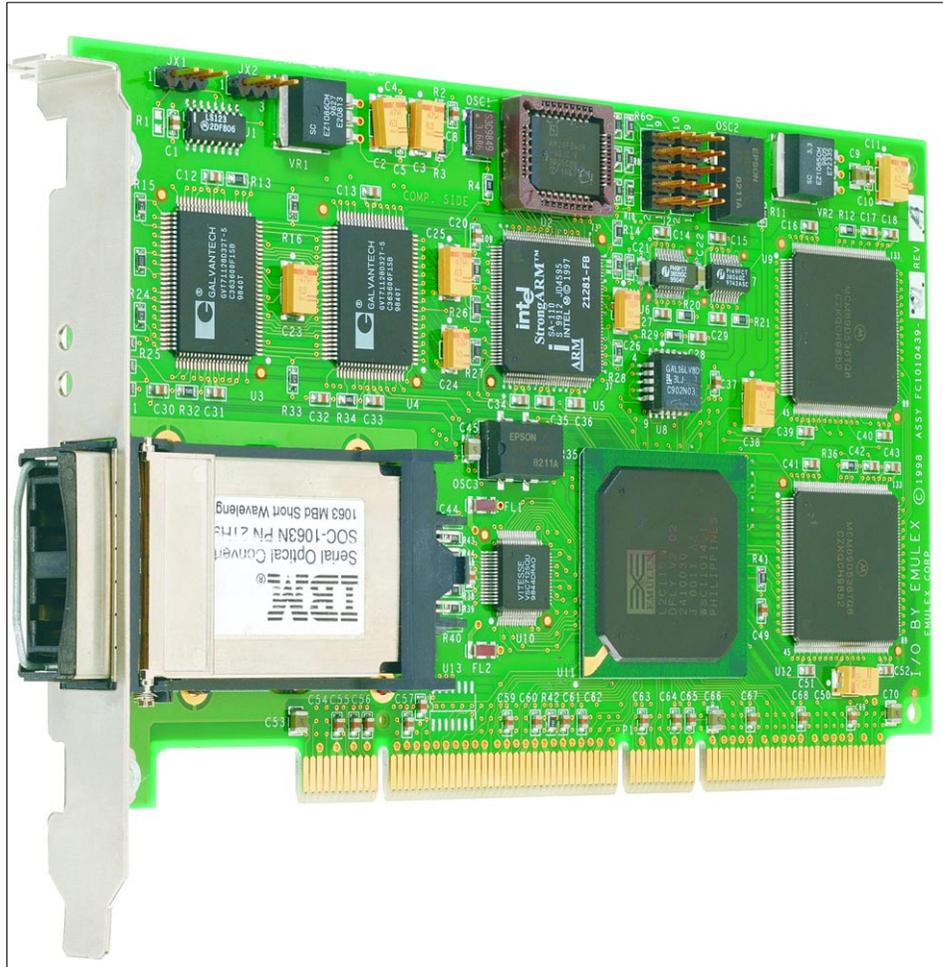


Figure 3-1 Emulex LP8000 adapter

- ▶ JNI FCI-1063 - PCI adapter: This adapter requires the latest driver from JNI. There is no firmware for this adapter. The driver download file is available at the following URL:
<http://www.jni.com/Drivers/drivers.3.cfm?ID=8>
- ▶ JNI FC64-1063 - Sbus adapter: This adapter requires the latest driver from JNI. There is no firmware for this adapter, which is shown in Figure 3-2. The driver download file is available at the following URL:
<http://www.jni.com/Drivers/drivers.3.cfm?ID=10>



Figure 3-2 JMI FC64-1063 adapter

- ▶ JMI FCE-6460 - PCI adapter: This is the latest JMI adapter and is currently being tested by IBM. This adapter requires the latest driver from JMI (currently v5.1). There is no firmware for this adapter. Unlike other adapters, this driver requires that the WWN of the targets be defined prior to it probing the fabric for them.
- ▶ QLogic QLA2200F - PCI adapter: This adapter requires the latest firmware, and driver from QLogic. Figure 3-3 is a picture of the this adapter. All downloads including documentation, are available at the following URL:
http://www.qlogic.com/support/product_resources.asp?id=113



Figure 3-3 QLogic QLA2200F adapter

Note: QLogic QLA2200F requires driver V3.07 for support.

3.2.1 Tuning HBA configuration files

With the Solaris operating system, there are a number of parameters which must be tuned for proper operation with Fibre Channel usage. These parameters are FC HBA dependent, but reside in the standard Solaris kernel. These parameters are outlined in Table 3-2 for the adapters which have been proofed by IBM.

For other adapters, check with your Sun representative or the adapter vendor.

Table 3-2 Solaris kernel parameter settings required

kernel parameter	JNI adapters settings	Emulex adapter settings	QLogic adapter settings
sd_io_time	0x78	0x78	0x78
sd_retry_count	5	5	5
maxphys	maxphys=8388608	maxphys=8388608	maxphys=8388608
sd_max_throttle	256/highest LUN count per controller	N/A	N/A

In addition, Fibre Channel host bus adapter configuration files are tunable for host system reliability and performance.

<http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm>

See Table 3-3 for the recommended configuration settings for the host-bus-adapter for a JNI FC64-1063 and a JNI FCI-1063.

Table 3-3 Recommended JNI HBA settings

Parameters	Recommended settings
fca_nport	0: Default. Initializes on a loop 1: Recommended for fabric. Initializes as an N_Port.
public loop	0: Default. Recommended. Initializes according to which fca_nport is set as disabled.
ip_disable	0: Default. IP side of the driver is enabled. 1: Recommended for fabric. The IP side of the adapters is completely disabled.
failover	See the following Web site for details on how to set the parameter setting: http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm
busy_retry_delay	500: Recommended. Delay between retries after device returns a busy response for a command.
scsi_probe_delay	5000: Recommended. Delay before SCSI probes are allowed during startup.

In Table 3-4 we show the recommended settings for the JNI FCE-6460.

Table 3-4 JNI FCE-6460 settings

Parameters	Recommended
FcLoopEnabled	1: Default. Initializes on a FC loop 0: Turns off loop protocol for fabric
FcFabricEnabled	0: Default Turns off N_Port initialization for FC loop 1: For fabric. Initializes as an N_Port
jnic146xn-def_hba_binding="null"	Set to ="null" this is an added parameter that must be added for each adapter in the system (where 'n' is the adapter number)

Parameters	Recommended
jnic146xn-target1_wwnn	Set to equal the WWN of the target devices port adapter(s) that this HBA is able to access (where 'n' is the HBAs number in the system)
jnic146xn-targetn_lunn_hba	Set to equal "jnic146xn" (where 'n' is the adapter/target/lun numbers). This entry must be made for each device the adapter is to have access to.
jnic146xn-targetn_lunn_throttle	Equal to 60 (where 'n' is the adapter/target/lun number)

See Table 3-5 for recommended host bus adapter configuration file parameters for an Emulex LP8000 adapter.

Table 3-5 Recommended Emulex LP8000 HBA settings.

Parameters	Recommended settings
automap	1: Default. Automatically assigns SCSI IDs to FCP targets.
fcp-on	1: Default. Turn on FCP.
lun-queue-depth	16: Recommended when there are less than 17 LUNs per adapter. Set value = 256/(Total LUNs per adapter) when there are more than 16 LUNs per adapter. If your configuration includes more than one LP8000 adapter per server, calculate the lun-queue-depth value using the adapter with the most LUNs attached.
no-device-delay	15: Recommended. Delay to failback and I/O.
network-on	0: Default. Recommended for fabric. Do not turn on IP networking 1: Turn on IP networking.
scan-down	2: Recommended. Use inverted ALPA map and cause target assignment in private loop.
topology	2: Recommended for fabric. Point-to-point mode only. 4: Recommended for non-fabric. Arbitrated loop mode only.
zone-rscn	0: Default 1: Recommended for fabric; check name server for RSCNs.

In Table 3-6 we show the recommended configuration file parameters for the host bus adapters for the QLogic QLA2200F adapter.

Table 3-6 Recommended QLA2200F HBA settings

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-enable-adapter-hard-loop	0
hba0-adapter-hard-loop-l	0
hba0-enable-64bit-addressing	0
hba0-enable-LIP-reset	0
hba0-enable-LIP-full-login	1
hba0-enable-target-reset	0: Non-clustered 1: Clustered
hba0-reset-delay	5
hba0-port-down-retry-count	30
hba0-link-down-error	1
hba0-loop-down-timeout	60
hba0-connection-options	1: Fabric connection 2: Direct connection
hba0-device-configuration-mode	1
hba0-fc-tape	0
hba0-command-completion-option	1

For other vendors, you must contact your Sun representative for assistance.

3.3 Sun supported Fibre Channel switches and hubs

In this section we discuss the different Fibre Channel switches which are available for use with the Sun platforms in the SAN environment. Again, Sun has partnered with a number of vendors to create quite a list of supported products which we have listed here with their vendor's URLs:

- ▶ StorEdge Network FC Switch-8, and 16 Sun Microsystems
http://www.sun.com/storage/fcswitch/san_8_16_datasheet.pdf
- ▶ Silkworm Series, Brocade Communications Systems Inc.
<http://www.brocade.com/products/>
- ▶ ED5000 (2032-001), ED6064 (2032-064), ES3016 (2032-016), ES3032 (2032-032), McDATA Corporation
www.mcdata.com/IBM/suptdev/
- ▶ FC/9000 Series, INRANGE Corporation
<http://www.inrange.com/products/san/fc9000>
- ▶ SANbox, SANbox2 Series, QLogic Corporation (Previously Ancor Communications Inc.)
http://www.qlogic.com/products/fc_san_switchs.asp
- ▶ 8100 switch, 1000/2000 hubs, Vixel Corporation
<http://www.vixel.com/products/>

IBM proofed switches with Sun systems

IBM has proofed the Sun systems on the IBM 2109, INRANGE 9000, McDATA ED5000 series, and will be testing it with the McDATA ED6064. The following brief descriptions provide a brief introduction to them. For greater detail, contact the specific switch vendor.

3.3.1 IBM 2109

Current members of the IBM 2109 family of switches are described in Table 3-7.

Table 3-7 The IBM 2109 family

Model	Ports	Speed	Connectors	Rack space
IBM TotalStorage SAN Switch S08	8	1 Gb/s	SC	1U
IBM TotalStorage SAN Switch S16	16	1 Gb/s	SC	2U
IBM TotalStorage SAN Switch F16	16	1 or 2 Gb/s	LC	1U

The 2109-S08, and 2109-S16 are shown in Figure 3-4. At the time of writing the 2109-F16 is under test with the Sun systems.



Figure 3-4 IBM 2109-S08 (top) and S16

Features and limitations

Special features of the IBM 2109 switch family are:

- ▶ Switch domain ids 1-239 are supported.
- ▶ The IBM 2109 family of switches support both fabric-attached and loop-attached devices on any port.
- ▶ Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on customer requirements

3.3.2 INRANGE FC/9000

The INRANGE directors currently sold by IBM (the IBM 2042 family) are described in Table 3-8.

Table 3-8 INRANGE directors

Model	Ports	Speed	Connectors	Rack space
2042-001 INRANGE FC/9000 Fibre Channel Director	64	1 Gb/s	SC	20U
2042-128 INRANGE FC/9000-128 Fibre Channel Director	128	1 Gb/s	SC	40U

In addition, IBM also sells the INRANGE FC/9000 40U cabinet (2042-C40). The cabinet is able to hold one 128-port switch or two 64-port switches. In Figure 3-5 we show the INRANGE FC/9000.



Figure 3-5 INRANGE FC/9000 switch

Features and limitations

Special features of the INRANGE FC/9000 are:

- ▶ The INRANGE FC/9000 family of switches support both fabric-attached and loop-attached devices on any port.
- ▶ Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on customer requirements.

3.3.3 McDATA

The McDATA components currently sold by IBM (the IBM 2031 and 2032 families) are described in Table 3-9.

Table 3-9 *McDATA components*

Model	Ports	Speed	Connectors	Rack space
2031-016 McDATA ES-3016 Fabric Switch	16	1 Gb/s	LC	1U
2031-032 McDATA ES-3032 Fabric Switch	32	1 Gb/s	LC	1.5U
2031-L00 McDATA ES-1000 Loop Switch	8+1	1 Gb/s	SC	1U
2032-001 McDATA ED-5000 Director with 32 ports	32	1 Gb/s	SC	18U
2032-064 McDATA ED-6064 Enterprise Fibre Channel Director with 64 ports	64	1 Gb/s	LC	9U

Note: McDATA ED5000 is not supported with QLogic QLA2100F FC HBA.

In addition, IBM also sells the McDATA FC-512 Cabinet (2032-C36), which can house up to four ED-6064 rack units. Both of these are shown in Figure 3-6.



Figure 3-6 McDATA cabinet (left), and 6064 rack mountable switch

Features and limitations

Special features and limitation of the McDATA switches are:

- ▶ The only McDATA switch that supports loop attached devices (tape drives, for example) is the McDATA ES-1000 Loop Switch.
- ▶ The ports in the 2032-064 can be added in cards. Each card contains four ports. There are three different cards available:
 - Four shortwave (multimode) ports.
 - Four longwave (singlemode) ports.
 - Three shortwave ports and one longwave port.
- ▶ On the 2031-016 and 2031-032 switches, any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC.
- ▶ The 2031-L00 switch has two different types of ports:
 - Eight switching Hub ports that can only be configured with shortwave (multimode) GBIC (or copper interface).
 - One bridge port, for connecting the loop switch to the fabric, that can be configured with a shortwave (multimode) or longwave (singlemode) GBIC.
- ▶ The McDATA switches support fabrics of up to 31 switches only.

3.3.4 Fibre Channel switch installation preparation

If you are installing the Sun system on a new, unused Fibre Channel switch, you should set up the fabric as a single zone, which allows you to familiarize yourself with the SAN environment before expanding it. If, however, you are adding the Sun to an already existing fabric, you may need to configure your switch to work with multiple hosts and different vendors and operating systems. For this, more planning is required.

Using zones with the switches

When you are sharing a switch with other hosts systems in a heterogeneous environment, you will need to know what the requirements of all the host vendors are with respect to sharing. Most vendors will require their host be configured in their own separate zones. This is to prevent the risk of any host from impacting the others, or their data on the fabric. This has been proven to be a good practice with all heterogeneous environments. The type of zoning you can use will vary from switch vendor to vendor; as does the methods that they use to implement the zoning types on their switches. Review closely the details that the vendor provides to ensure that you understand the protection that your zone will be providing you.

Using cascaded switches

Sun systems can be used with cascaded fabric environments as long as they meet the specifications of the switch vendors. Again, we strongly recommend that you review your planned configuration with the appropriate representatives to ensure compatibility, and functionality. We especially recommend care when considering mixed (heterogeneous) switched environments. Heterogeneous fabrics have not been fully tested by IBM with Sun platforms.

3.4 Supported storage on the Sun fabric

The Sun platform supports a great number of storage devices on its fabrics. These are both Sun's own StorEdge line, as well as a high number of third party storage devices. For a complete list of supported devices contact your Sun representative.

For our purpose, we list the IBM supported devices for the Sun platforms in Table 3-10 and Table 3-11.

Table 3-10 IBM Fibre Channel attached disk storage systems

Storage system	Method of attachment	Operating systems
IBM 3542/3560 - FAStT200 Storage Server with up to five EXP500 disk expansion unit	Attach by both fibre switch, or FC-AL.	Supported with Solaris 2.6, 7, and 8.
IBM 3552/3560 - FAStT500 with up to 22 EXP500 disk expansion units	Attach by both fibre switch, or FC-AL.	Supported with Solaris 2.6, 7, and 8.
IBM 1742/3560 - FAStT700 with up to 22 EXP500 disk expansion units	Attach by both fibre switch, or FC-AL.	Supported with Solaris 2.6, 7, and 8.
IBM 2106 Modular Storage Server (MSS)	Attach by both fibre switch, or FC-AL.	Supported with Solaris 2.6, 7, and 8.
IBM 7133 - Serial Disk System, with 7140 SAN controller 160	Attach by both fibre switch, or FC-AL.	Supported with Solaris 2.6, 7, and 8.
IBM 2105 Enterprise Storage Server (ESS)	Attach by both fibre switch, or FC-AL.	Supported with Solaris 2.6, 7, and 8.

Table 3-11 IBM Fibre Channel attached tape storage systems

Storage system	Method of Attachment	Operating systems
IBM 3584 Ultra Scalable Tape Library Server (LTO)	FC-AL only	Supported with Solaris 2.6, 7, and 8.
IBM 3590E tape drives with direct fibre interface	FC-AL only	Supported with Solaris 2.6, 7, and 8.

3.5 Installation planning

Before starting the installation, it is important to plan what you intend to do. This includes deciding what will be connected and where the equipment will be located. The example installation which follows will be small, but your installation may be quite large. Items to consider are the location of your system, the storage that is being connected, and the switch(es) which you are connecting them to.

It is best to map out the entire layout prior to starting any work. You will want to know the power requirements, along with distance, and routing of all fiber optic cables you will use.

3.5.1 Cabling requirements and considerations

Cables must be purchased to connect these all together; or possibly an installer will need to be called, especially if running between floors of a building. Ready made up cable lengths are available from some vendors for most common installations.

With fibre cables, a point to remember is that attenuation is your worst enemy; it can cause your entire network to be highly unstable, or not to function at all. Watching your number of connections, ensuring clean ends, and proper seating, is very critical to your success. Do not try to stretch the maximum distances of the signal.

In Table 3-12 we show the cable types and their maximum limits.

Table 3-12 Common Fibre Channel cabling configurations

GBIC type	Cable type	Distance (1 Gb/s)	Distance (2 Gb/s)
longwave	9 micron, singlemode	10 km	10 km
shortwave	50 micron, multimode	500 m	300 m
shortwave	62.5 micron, multimode	300 m	150 m

Ensure that any cables you buy or have made are constructed with the correct connector ends to match with your switch ports and FC HBA (either SC, or LC).

3.6 Installation procedure

In this section we run through a simple fabric installation and show the success of the work. We are using a Sun 450 system with JNI's FCE-6460 PCI FC HBAs installed, connected to an IBM ESS using one FC host adapter. We have LUNs configured on the ESS for access by the Sun system. We are also using Solaris V2.6 to perform our installation. Any special instructions which may differ from our performed steps for the other supported releases of the operating system are noted as we go.

To start with it is assumed that this not a new system installation and that the Fibre Channel we are installing is an add-on to the platform hardware. Also the ESS has been setup and configured to provide the LUNs to our system once we are ready to provide it with the correct WWNs. The system is powered up and running the operating system of choice. It is also assumed that the system has been configured for LAN and TCP/IP support. Now we are ready to start installing our fabric.

3.6.1 Check that the system is at a proper level

Before we start our installation, we first make some checks to ensure that the system is running the version of operating system we desire, the correct patches, and does not already have the driver loaded. This we do by running the following commands.

1. To check the operating system level, we run:

```
# uname -a
SunOS sol-e 5.6 Generic_105181-29 sun4u sparc SUNW,Ultra-4
```

2. To ensure that we have the necessary patches loaded, we use the following command with the appropriate patch number:

```
# showrev -p | grep 105181 (perform this command for each patch number.)
[Result goes here]
# showrev -p | grep 105356
[Result goes here]
# showrev -p | grep 105580
[Result goes here]
# showrev -p | grep 105600 (made obsolete in our case by 105181-29)
[Result goes here]
```

3. To check for the driver's existence, we run:

```
# pkginfo | grep -i jni
```

At this point the system is ready for the adapter to be installed.

3.6.2 Installing the FC HBA adapter card

In this section, we explain the steps we performed to install the JNI FCE-6460 PCI adapter card. These steps are the same for installing all PCI based adapters.

1. Turn off and unplug the computer.
2. Remove the computer case.
3. Remove the blank panel from an empty PCI bus slot.
4. Insert the host adapter board into the empty PCI bus slot. Press firmly until seated.
5. Secure the mounting bracket for the adapter to the case with the panel screw.
6. Replace the computer case by tightening the screws on the case or use the clamp to secure the cover.

3.6.3 Switch setup

In this section, we explain the steps we took to install the switch, and prepare it to be connected.

1. Using the switch's *local control panel*, we set the switch to its default settings and password; and also set the TCP/IP address (9.43.231.91) as it will be needed to be able to access it using the remote telnet login to complete the configuration as needed. This IP address can also be used to access the switch via the Web based GUI session.
2. Now we telnet to the switch with the IP address and login using the default login, and password (login=admin, password= password).

```
# telnet 9.43.231.91
Trying...
Connected to 9.43.231.91.
Escape character is '^]'.
Fabric OS (tm) Release a2.4.1c
login: osl
Password:
osv12109a:osl>
```

3. We confirmed these settings by running the following **configShow** command.

```
osv12109a:osl> configShow
Ethernet addresses: 0:60:69:12:3f:fc
Nvram datadata: fei(0,0)host:/usr/switch/firmware e=9.43.227.19:fffffc00
g=9.43.
224.1 u=user tn=osv12109a
diag.postDisable:      0
fabric.domain:         1
fabric.ops.BBCredit:   16
fabric.ops.E_D_TOV:    2000
fabric.ops.R_A_TOV:    10000
fabric.ops.dataFieldSize: 2112
fabric.ops.mode.SeqSwitching: 0
fabric.ops.mode.fcpProbeDisable: 0
fabric.ops.mode.isolate: 0
fabric.ops.mode.longDistance: 0
fabric.ops.mode.noClassF: 0
fabric.ops.mode.pidFormat: 0
fabric.ops.mode.sync: 0
fabric.ops.mode.tachyonCompat: 0
fabric.ops.mode.unicastOnly: 0
fabric.ops.mode.useCsCtl: 0
fabric.ops.mode.vcEncode: 0
fabric.ops.vc.class.2: 2
fabric.ops.vc.class.3: 3
fabric.ops.vc.config: 0xc0
fabric.ops.vc.linkCtrl: 0
fabric.ops.vc.multicast: 7
fc4.fcIp.address:      0.0.0.0
fc4.fcIp.mask:         0.0.0.0
fcAL.fanFrameDisable: 0
```

```

fcAL.openSendCLS:      0
fcAL.useAltBBCredit:  0
lcdContrast:          128
lcdContrast.orange:   208
licenseKey:           zdeSSR9zy2efTTY;SbRcyRQSDcTcSp0y
route.delayReroute:   0
route.embeddedPortBcast: 1
route.stickyRoutes:   1
rpc.rapid:            1
rpc.rstatd:           0
rpc.rusersd:          0
shell.delete:         0
shell.quiet:          0
shell.timeout:        0
syslog.errDisplayLevel: 5
syslog.errsOnConsole: 1
thresh.alarmsfilter:  1
thresh.thad:          1
xrelativeModeDisable: 0
Displaying enabled/disabled portlog events
Type <CR> to continue, Q<CR> to stop: q

```

4. We now need to ensure we have the proper switch firmware. This we do by running the **version** command on the switch telnet session.

```

osv12109a:os1> version
Kernel:      5.3.1
Fabric OS:   a2.2.1a
Made on:     Tue May 8 15:44:06 PDT 2001
Flash:      Tue May 8 15:45:24 PDT 2001
BootProm:   Thu Jun 17 15:20:39 PDT 1999

```

5. We now need to connect the cables ensuring that we are careful not to touch the ends and making sure they are free of any dirt or dust. We seat them securely into the GLM of the FC HBA, and the GBIC of the switch port. By issuing the **switchShow** command, we can see these switch details:

```

osv12109a:os1> switchShow
switchName:   SNJ2109DD2
switchType:   2.4
switchState:  Online
switchRole:   Principal
switchDomain:  1
switchId:     fffc01
switchWwn:    10:00:00:60:69:12:3f:fc
switchBeacon: OFF
port 0: sw No_Light
port 1: sw No_Light

```

```
port 2: sw No_Light
port 3: sw No_Light
port 4: -- No_Module
port 5: -- No_Module
port 6: -- No_Module
port 7: -- No_Module
port 8: -- No_Module
port 9: -- No_Module
port 10: -- No_Module
port 11: -- No_Module
port 12: sw No_Light
port 13: sw No_Light
port 14: sw No_Light
port 15: sw No_Light
```

3.6.4 Storage setup

In this section, we explain the steps we took to prepare and connect the storage device which, in this example, is an ESS.

1. We need to ensure that our storage device is ready to be connected to our Sun system. In our example we are using a group of LUNs from an IBM ESS. So, we must now define the FC HBA to the ESS using the StorWatch manager. As we are sharing the ESS between multiple hosts and the Sun, the ESS ports will be configured for “Access Restricted”, and specific LUN assignments will be used. To do this we must map the desired LUNs to the FC HBA’s WWN. We obtain the “Node WWN” of the FC HBA from the **infodev** command we ran in step 3 above.
2. We now need to connect the cables from the ESS to the switch, again being careful not to contaminate them with dirt or oil from our hands. We seat them securely into the GLM of the ESS adapter, and the GBIC of the switch port. Another check using **switchShow** gives us the following:

```
osv12109a:os1> switchShow
switchName:   osv12109a
switchType:   2.4
switchState:  Online
switchRole:   Principal
switchDomain:  1
switchId:     fffc01
switchWwn:    10:00:00:60:69:12:3f:fc
switchBeacon: OFF
port 0: sw No_Light
port 1: sw No_Light
port 2: id No_Light
port 3: sw No_Light
port 4: -- No_Module
port 5: -- No_Module
```

```
port 6: id Online          F-Port 50:05:07:63:00:c0:8b:2b
port 7: -- No_Module
port 8: -- No_Module
port 9: -- No_Module
port 10: sw No_Light
port 11: sw No_Light
port 12: sw No_Light
port 13: sw No_Light
port 14: sw No_Light
port 15: id No_Light
```

3.6.5 Downloading the Fibre Channel adapter drivers

We show how to download the JNI, Emulex and QLogic drivers in the following topics.

Downloading the JNI adapter driver

In this section we explain how to download the JNI 6460 adapter driver:

1. Plug in and restart your host system, and login as “root”.
2. Enter the following to create a temporary directory.

```
# mkdir /jnitemp
```
3. Move to the temporary directory with the following command.

```
# cd /jnitemp
```
4. Go to the following Web site:
<http://www.jni.com>
5. From the navigation menu at the top of the page, click **Drivers**.
6. From the menu for Locate Driver by Product, click **FCE-6460**.
7. From the menu for **FCE-6460**, find the section for **Solaris** and copy or print the **readme.txt** file for reference while installing.
8. Now, from the menu for **FCE-6460**, and the section for Solaris select **JNIC146x.pkg**, and click to download.
9. In the Save As dialog box, we directed it to our temporary folder **/jnitemp**, and left the filename as **JNIC146x**, and clicked **Save**. When you click Save, a window displays that shows the progress of the download.
10. When the download completes, click **Close**.

Note: If you are installing a different JNI adapter, the process is the same, except that you would pick the appropriate adapters's driver file to download from the Web page.

Downloading the Emulex adapter driver

Following is the step-by-step process to follow to download the Emulex driver to your system for installation:

1. Plug in and restart your host system, and login as “root”.
2. Enter the following to create a temporary directory.

```
# mkdir /emulextemp
```
3. Move to the temporary directory with the following command.

```
# cd /emulextemp
```
4. Go to the following Web site:
<http://www.emulex.com/ts/dds.html>
5. From the **Models** selection box, select **LP8000** (or whatever your adapter type may be). When you click on your model you will be automatically taken to the relevant documentation and driver page for downloading.
6. Scroll down to the section titled **Drivers for Solaris** and select the **PCI Version** download.
7. From the Solaris Driver Version screen select either the **Intel** or **Sparc** download version depending on your platform you will be running it on.
8. From the **Open it/Save to disk** menu select **Save to disk**, and save the file to your **/emulextemp** directory.
9. The file will save as a tar file, and will need to be extracted by doing the following:

```
# tar xvf lpfc-sparc.tar
```

Downloading the QLogic adapter driver

The following is the step-by-step process to follow to download the QLogic driver to your system for installation:

1. Plug in and restart your host system, and login as “root”.
2. Enter the following to create a temporary directory.

```
# mkdir /qlogictemp
```
3. Move to the temporary directory with the following command.

```
# cd /qlogictemp
```

4. Go to the following Web site:
http://www.qlogic.com/support/home_support.asp
5. From the **HBA** selection box, select the product **QLA2200F** (or whatever your adapter type may be). When you click on your model you will be automatically taken to the relevant documentation and driver page for downloading.
6. Scroll down to the section titled **Drivers** and select **Solaris**.
7. From the **Description** section select **Installation** and download the readme file to your system for reference.
8. From the **Download** section, select **download** to download the driver to your system and save it to your `/qlogictemp` directory.
9. The file will save as a compressed tar file, and will need to be uncompressed, and untarred into the directory for installation.

```
# uncompress q1a2200.Z
# tar xvf q1a2200.tar
```

Note: If we had downloaded the file to another platform, we would need to move it to our system for installation, normally with something like “ftp”.

Additionally, some vendors package their driver files in tar files, and may also compress them. These need to be uncompressed and untarred prior to installation. Commands for these steps are:

```
# uncompress [filename.tar.Z]
# tar xvf [filename.tar]
```

The result is the pure driver files you will need to install (possibly in a sub-directory structure).

We are now ready to install the driver on our system.

3.6.6 Installing the Fibre Channel adapter driver(s)

In this section we are installing the FC HBA driver for the adapter. The steps here are for the JNI 6460, but are generally the same for all the driver vendors. To install the Fibre Channel adapter drivers, we performed the following steps:

1. If you have not already done so, login as root.
2. Change directory to the `jnitemp` (or whatever temp directory you used for your adapter’s driver).
3. Create a subdirectory tree under `/opt/Drivers/` called `JNI/JNIC146x/` (or whatever name is appropriate for your adapter’s driver):

```
# mkdir /opt/Drivers/JNI
#
# mkdir /opt/Drivers/JNI/JNIC146x
#
```

4. Copy the device driver file to the newly created `/opt/Drivers/JNI/JNIC146x` directory.

```
# cp /jnitemp/JNIC146x.pkg /opt/Drivers/JNI/JNIC146x/
```

5. Now we install the driver with the following command, and answer the script questions as needed:

```
# pkgadd -d /opt/Drivers/JNI/JNIC146x/JNIC146x.pkg
```

Note: An installation script displays that prompts you to answer a number of questions. For each question, enter the appropriate response, or press Enter to each question to accept the default setting.

The following packages are available:

```
1  JNIC146x      JNI Fibre Channel SCSI HBA Driver
                   (sparc) v5.1 (01110800)
2  JNIsnia      JNI SNIA Fibre Channel HBA LIBRARY (Solaris)
                   (sparc) v1.0.0.0.0.b.011106-10
```

Select package(s) you wish to process (or 'all' to process all packages).
(default: all) [?,??,q]:

```
Processing package instance <JNIC146x> from
</opt/Drivers/JNI/FCE-6460/v5.1/JNIC146x.pkg>
```

```
JNI Fibre Channel SCSI HBA Driver
(sparc) v5.1 (01110800)
Copyright (c) 1999-2001 JNI Corp., All Rights Reserved.
## Processing package information.
## Processing system information.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.
```

This package contains scripts which will be executed with super-user permission during the process of installing this package.

Do you want to continue with the installation of <JNIC146x> [y,n,?] y

```
Installing JNI Fibre Channel SCSI HBA Driver as <JNIC146x>
```

```

## Installing part 1 of 1.
/kernel/drv/jnic146x
/kernel/drv/jnic146x.conf
/kernel/drv/sparcv9/jnic146x
[ verifying class <none> ]
## Executing postinstall script.

Installation of <JNIC146x> was successful.

Processing package instance <JNISnia> from
</opt/Drivers/JNI/FCE-6460/v5.1/JNIC146x.pkg>

JNI SNIA Fibre Channel HBA LIBRARY (Solaris)
(sparc) v1.0.0.0.0.b.011106-10
Copyright 2001 JNI Corp., All Rights Reserved.
## Processing package information.
## Processing system information.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <JNISnia> [y,n,?] y

Installing JNI SNIA Fibre Channel HBA LIBRARY (Solaris) as <JNISnia>

## Installing part 1 of 1.
/opt/JNISnia/Solaris/Common/32bit/HbaCommonLib.so
/opt/JNISnia/Solaris/Common/64bit/HbaCommonLib.so
/opt/JNISnia/Solaris/Common/Sample/JniHbaSample.c
/opt/JNISnia/Solaris/Common/Sample/Makefile
/opt/JNISnia/Solaris/Common/Sample/hbaapi.h
/opt/JNISnia/Solaris/Common/Sample/testlib32
/opt/JNISnia/Solaris/Common/Sample/testlib64
/opt/JNISnia/Solaris/Jni/32bit/JniHbaLib.so
/opt/JNISnia/Solaris/Jni/64bit/JniHbaLib.so
/opt/JNISnia/Solaris/Jni/Sample/JniHbaSample.c
/opt/JNISnia/Solaris/Jni/Sample/Makefile
/opt/JNISnia/Solaris/Jni/Sample/hbaapi.h
/opt/JNISnia/Solaris/Jni/Sample/testlib32
/opt/JNISnia/Solaris/Jni/Sample/testlib64
/opt/JNISnia/Solaris/SNIA_readme.txt
/opt/JNISnia/Solaris/hba.conf
/opt/JNISnia/Solaris/hbaapi.h
[ verifying class <none> ]
## Executing postinstall script.
----- Checking for previously installed (SNIA) /etc/hba.conf

```

```

----- Existing hba.conf found - Update in progress
----- Existing hba.conf saved as - hba.conf.bak.1204.080938 -
----- hba.conf updated
----- hba.conf ready
Installation of <JNIsnia> was successful.

```

Note: With some drivers you will be prompted to rebuild the manpages to add the drivers manpage(s) to the system library. This rebuild can take up to as much as ten minutes. If you do not want to do this now type **n**. You can run this command later.

6. To check that the installation of the driver was truly successful, we enter the following command. Substitute your driver's name for `jni` as appropriate.

```

# pkginfo | grep -i jni
system      JNIC146x      JNI Fibre Channel SCSI HBA Driver
application JNIsnia       JNI SNIA Fibre Channel HBA LIBRARY (Solaris)

```

Before we reboot the system we must configure the necessary parameters for our adapter. It is strongly recommended that you do this for all adapters, prior to restarting the system.

3.6.7 Configuring the host device and adapter drivers

At this point we needed to make the necessary changes to the device driver configuration file, and the FC HBA driver configuration file to support the adapter as desired, and to allow access to the devices. To do this, we followed these steps:

1. Change to the directory by typing:

```
# cd /kernel/drv
```

2. Back up the `sd.conf` file in this subdirectory:

```
# cp sd.conf sd.orig.conf
```

3. Edit the `sd.conf` file to add support for the target and LUN pairs that are configured on the host system. The following is what we added for our example:

```

# vi sd.conf
sd.conf" [Read only] 587 lines, 21887 characters
# Copyright (c) 1992, by Sun Microsystems, Inc.
#
#ident"@(#)sd.conf1.893/05/03 SMI"

name="sd" class="scsi" class_prop="atapi"
    target=0 lun=0;
name="sd" class="scsi" class_prop="atapi"

```

```

    target=1 lun=0;
name="sd" class="scsi" class_prop="atapi"
    target=2 lun=0;
name="sd" class="scsi" class_prop="atapi"
    target=3 lun=0;
name="sd" class="scsi"
    target=4 lun=0;
name="sd" class="scsi"
    target=5 lun=0;
name="sd" class="scsi"
    target=6 lun=0;
name="sd" class="scsi"
    target=8 lun=0;
name="sd" class="scsi"
    target=9 lun=0;
name="sd" class="scsi"
    target=10 lun=0;
name="sd" class="scsi"
    target=11 lun=0;
name="sd" class="scsi"
    target=12 lun=0;
name="sd" class="scsi"
    target=13 lun=0;
name="sd" class="scsi"
    target=14 lun=0;
name="sd" class="scsi"
    target=15 lun=0;
#
#Lines added for ESS on JNI 6460 controller:
name="sd" class="scsi"
    target=0 lun=1;
name="sd" class="scsi"
    target=0 lun=2;
name="sd" class="scsi"
    target=0 lun=3;
name="sd" class="scsi"
    target=0 lun=4;
name="sd" class="scsi"
    target=0 lun=5;
name="sd" class="scsi"
    target=0 lun=6;
name="sd" class="scsi"
    target=0 lun=7;
name="sd" class="scsi"
    target=0 lun=8;
:wq

```

Note: Do not add duplicate target and LUN pairs.

4. Now we edit the `driver.conf` file for our driver to have the specific parameter settings that are needed to work properly with our Sun system. In our example we edited the `jnic146x.conf` file. Refer to the vendor's documentation for your adapters settings.

Note: In Table 3-3 on page 66, Table 3-5 on page 67, and Table 3-6 on page 68 we show the recommended parameters for the IBM tested adapters.

```
# vi jnic146x.conf (and add the following:)

FcFabricEnabled = 0;
FcFabricEnabled = 1;
jnic146x0-def_hba_binding="null"; (for the first adapter)
jnic146x1-target1_wwnn="$5005076300c08b2b"; (from the ESS host port we will
                                         access the devices through)

jnic146x1-target0_lun1_hba="jnic146x1";
jnic146x1-target0_lun1_throttle=60;
jnic146x1-target0_lun2_hba="jnic146x1";
jnic146x1-target0_lun2_throttle=60;
jnic146x1-target0_lun3_hba="jnic146x1";
jnic146x1-target0_lun3_throttle=60;
jnic146x1-target0_lun4_hba="jnic146x1";
jnic146x1-target0_lun4_throttle=60;
jnic146x1-target0_lun5_hba="jnic146x1";
jnic146x1-target0_lun5_throttle=60;
jnic146x1-target0_lun6_hba="jnic146x1";
jnic146x1-target0_lun6_throttle=60;
jnic146x1-target0_lun7_hba="jnic146x1";
jnic146x1-target0_lun7_throttle=60;
:wq
```

5. We now have the drivers ready to be reconfigured into the kernel. But before we do this, we want to connect the system to our switch that we will use.
6. Now we are ready to shut down and restart the system. We do this with the following command:

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

This will cause the kernel to reconfigure during the reboot.

7. When reboot completes, we should be able to see our devices by simply running the following command:

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
 0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /pci@1f,4000/scsi@3/sd@0,0
```

1. c0t2d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248> tre_SOL6
/pci@1f,4000/scsi@3/sd@2,0
 2. c0t3d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248> tre_sol7
/pci@1f,4000/scsi@3/sd@3,0
 3. c2t1d0 <IBM-2105F20-0.38 cyl 15256 alt 2 hd 30 sec 64>
/pci@1f,2000/JNI,FCR@1/sd@1,0
 4. c2t1d1 <IBM-2105F20-0.38 cyl 15256 alt 2 hd 30 sec 64>
/pci@1f,2000/JNI,FCR@1/sd@1,1
 5. c2t1d2 <IBM-2105F20-0.38 cyl 15256 alt 2 hd 30 sec 64>
/pci@1f,2000/JNI,FCR@1/sd@1,2
 6. c2t1d3 <IBM-2105F20-0.38 cyl 15256 alt 2 hd 30 sec 64>
/pci@1f,2000/JNI,FCR@1/sd@1,3
 7. c2t1d4 <IBM-2105F20-0.38 cyl 15256 alt 2 hd 30 sec 64>
/pci@1f,2000/JNI,FCR@1/sd@1,4
 8. c2t1d5 <IBM-2105F20-0.38 cyl 15256 alt 2 hd 30 sec 64>
/pci@1f,2000/JNI,FCR@1/sd@1,5
 9. c2t1d6 <IBM-2105F20-0.38 cyl 15256 alt 2 hd 30 sec 64>
/pci@1f,2000/JNI,FCR@1/sd@1,6
8. If devices do not get recognized with the JNI 6460 adapter, check for the following in the boot messages of the /var/messages file:

```
# pg /var/adm/messages
messages:Dec  4 08:09:10 trelane unix: jnic146x0: Hba: JNI,FCR Model:
FCE-6460-N
messages:Dec  4 08:09:10 trelane unix: jnic146x0: FCode: Version 3.6.2
[3b32]
messages:Dec  4 08:09:10 trelane unix: jnic146x0: IRQ: 1 Channel: 1
messages:Dec  4 08:09:10 trelane unix: jnic146x0: SubVend: 1242 Sub: 6460
Rev: 100 Vend: 1242 Dev: 1460
messages:Dec  4 08:09:10 trelane unix: jnic146x0: WWNN: 1000000173002f6e
WWPN: 2000000173002f6e
messages:Dec  4 08:09:10 trelane unix: jnic146x0: Auto-negotiating link
speed
messages:Dec  4 08:09:10 trelane unix: jnic146x0: Configured as Private
Loop port
messages:Dec  4 08:09:11 trelane unix: jnic146x0: JNIC v5.1 (01110800)
messages:Dec  4 08:09:11 trelane unix: jnic146x0: Copyright(c) 1995-2001
JNI Corp, All Rights Reserved.
messages:Dec  4 08:09:16 trelane unix: PCI-device: JNI,FCR@1, jnic146x #0
messages:Dec  4 08:09:16 trelane unix: JNI,FCR@1 is /pci@1f,2000/JNI,FCR@1
messages:Dec  4 08:09:21 trelane unix: jnic146x0: Link Failure event
detected
```

If you see these entries, then there is a good chance that you have something incorrect in your configuration file for the adapter. Normally a good candidate is the WWN entry for the device. Review your settings carefully. If you make changes to any of the **.conf** files, you must restart again with the **reboot -- -r** command.

9. A check which can be used with all adapters is the switch monitor tool. With the IBM 2109 we used a **switchShow** command to show whether the adapter is logged in to the switch or not. This is also useful in isolating the problem.

```
osv12109a:os1> switchShow
switchName:   osv12109a
switchType:   2.4
switchState:  Online
switchRole:   Principal
switchDomain:  1
switchId:     fffc01
switchWwn:    10:00:00:60:69:12:3f:fc
switchBeacon: OFF
port 0: sw No_Light
port 1: sw No_Light
port 2: id Online      F-Port  20:00:00:01:73:00:2f:6e
port 3: sw No_Light
port 4: -- No_Module
port 5: -- No_Module
port 6: id Online      F-Port  50:05:07:63:00:c0:8b:2b
port 7: -- No_Module
port 8: -- No_Module
port 9: -- No_Module
port 10: sw No_Light
port 11: sw No_Light
port 12: sw No_Light
port 13: sw No_Light
port 14: sw No_Light
port 15: id No_Light
```

For INRANGE switches, you should use the IN-VSN Enterprise Manager application to find the same information. For McDATA switches, you can find the information in the EFC Manager application. Both of these applications are graphical user interface (GUI) based for ease of use.

3.7 Device functionality checks

Now we are ready to checkout our disks to see if we can actually write to and read from them. We recommend the use of the **dd** or **rdcp** commands for this testing. An example of **dd** is shown here:

```
# dd if=/unix of=/dev/rdisk/c2t1d0s2 bs=1024 count=25
25+0 records in
25+0 records out
# dd if=/dev/rdisk/c2t1d0s2 of=/dev/null bs=1024 count=10
10+0 records in
10+0 records out
```

3.8 Specific limitations

The following is a list of some specific limitations which exist within some Sun environments and not others.

- ▶ Some of the newer FC HBAs use the above mentioned WWN entry for the device ports to control their fabric probe. This feature is not done by all adapters.
- ▶ The FC-AL environment cannot exceed 50 devices per loop.

For more details on these and other considerations that may exist in other environments, contact your Sun representative for help.



HP 9000 Fibre Channel implementation

In this chapter we discuss how to successfully set up a SAN fabric for the HP 9000 series platforms. We cover the step-by-step procedure from start to finish, and the supported products which work with the HP 9000 systems. We also point out where the latest firmware and device drivers can be obtained.

In this chapter, the following topics are discussed:

- ▶ “Introduction” on page 94
- ▶ “Fibre Channel host bus adapters” on page 94
- ▶ “SAN fabrics” on page 95
- ▶ “Storage” on page 99
- ▶ “Installation procedures” on page 100

4.1 Introduction

HP-UX is the operating system used on the HP 9000 systems.

There are three versions of HP-UX currently available, HP-UX v10.20, HP-UX v11.0, and HP-UX v11i. Fibre Channel operations are generally the same on these operating systems, and we discuss any known differences in their respective sections throughout this chapter.

In this chapter, we only cover a single-path environment, where each fibre device can be accessed via only one path. If you have a multi-path environment, you need to use either the HP-UX PV-Links feature or a storage system specific multi-path software, for example IBM Subsystem Device Driver for ESS.

4.2 Fibre Channel host bus adapters

In this topic we detail Fibre Channel host bus adapters (FC HBA) that can be implemented.

4.2.1 HP FC HBAs

There are three Fibre Channel adapters supporting switched fabrics that are currently available from HP. The adapters are described in Table 4-1. Each of these adapters supports only certain HP 9000 series models.

Table 4-1 Current HP FC HBAs with fabric support

Product number	Bus	Servers	HP-UX version
A5158A	PCI	Superdome, V-Class, N-Class, L-Class, A500, A400, rx9610, rx4610	11.0 and later
A6685A	HSC	K-Class	11i, 11.0, 10.20
A6684A	HSC eff	D-Class, R-Class, A180	11i, 11.0, 10.20

All of these adapters use SC connectors and have the speed of 1 Gb/s.

There are also several older adapters, that do not support switched fabrics. We do not cover those adapters.

You can find more information on these adapters and the servers supporting them at:

<http://techsolutions.hp.com/>

You can find the operating system patches and device drivers required for your adapter at:

<http://www.software.hp.com/>

4.2.2 Other FC HBAs

Other FC HBAs available for HP 9000 servers are listed in Table 4-2. These adapters are supported by their manufacturer. For more information about these adapters and their limitations, see the manufacturer's web site

<http://www.emulex.com/>

<http://www.jni.com/>

All of these adapters are PCI-bus adapters and can only be installed on PCI-bus based HP 9000 series systems.

Table 4-2 Other FC HBAs available for HP 9000 systems

Adapter	Ports	Connector	Speed	Notes
Emulex LP7000E	1	SC	1 Gb/s	
Emulex LP8000	1	SC	1 Gb/s	
Emulex LP8000DC	2	SC	1 Gb/s	
Emulex LP9000	1	SC	1 Gb/s	
Emulex LP9002L	1	LC	1 or 2 Gb/s	
Emulex LP9002DC	2	LC	1 or 2 Gb/s	
JNI FCE-3210	1	SC	1 Gb/s	
JNI FCE-6410	1	SC	1 Gb/s	
JNI FCE2-6410	2	SC	1 Gb/s	Partially supported
JNI FCE2-6412	2	SC	1 Gb/s	Partially supported

4.3 SAN fabrics

The HP 9000 series hosts support two Fibre Channel network topologies:

- ▶ Fibre Channel Arbitrated Loop (FC-AL)
- ▶ Switched fabric

The HP 9000 series systems have supported FC-AL loops for a long time. FC-AL has, however, severe limitations in performance and scalability. Because of those limitations we only cover switched fabric topology. However, we also discuss attaching storage devices supporting only FC-AL attachment to the fabrics.

We cover the products from the following three vendors:

- ▶ IBM 2109 family
- ▶ INRANGE
- ▶ McDATA

4.3.1 General information

There are two types of Gigabit Interface Converters (GBIC) and three types of fiber cable commonly used in Fibre Channel environments, as described in Table 4-3.

Table 4-3 Common Fibre Channel cabling configurations

GBIC type	Cable type	Distance (1 Gb/s)	Distance (2 Gb/s)
longwave	9 micron, singlemode	10 km	10 km
shortwave	50 micron, multimode	500 m	300 m
shortwave	62.5 micron, multimode	300 m	150 m

Shortwave GBICs and 50 micron cable are usually used to connect devices to the fabric, since most device adapters only support shortwave GBICs, as well as for short inter-switch links. Longwave GBICs and 9 micron cable are used when required for extended distance.

Attenuation

Fibre Channel is a high speed communication technology. Because of this, it is very sensitive to link attenuation. The attenuation is even more of a critical factor when, in the future, connection speeds rise from 1 Gb/s to 2 Gb/s and even higher.

To limit the effects of attenuation, we recommend making sure that:

- ▶ You have as few connections as possible on a single link.
- ▶ All connectors are clean.
- ▶ All connectors are properly seated.
- ▶ Only one type of fiber is used on a single link (never mix 50 micron and 62.5 micron fiber)

Zoning

We strongly recommend that you use zoning to:

- ▶ Enhance the security of your Fibre Channel fabric.
- ▶ Avoid interactions between different hosts.
- ▶ Secure storage devices that do not support LUN masking, like tape drives.

Since the implementation of zoning and even the terminology used varies between switch manufacturers, we do not discuss it in any detail.

Cascading

HP-UX supports cascading, but with only a maximum of two switches.

4.3.2 The IBM 2109 family

The current members of the IBM 2109 family of switches are described in Table 4-4.

Table 4-4 The IBM 2109 family

Model	Ports	Speed	Connectors	Rack space
IBM TotalStorage SAN Switch S08	8	1 Gb/s	SC	1U
IBM TotalStorage SAN Switch S16	16	1 Gb/s	SC	2U
IBM TotalStorage SAN Switch F16	16	1 or 2 Gb/s	LC	1U

Features and limitations

The IBM 2109 family of switches support both fabric-attached and loop-attached devices on any port, and switch domain ids 1-239 are supported.

Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on requirements.

For more information, see *IBM SAN Survival Guide Featuring the IBM 2109*, SG24-6127.

4.3.3 INRANGE

The INRANGE switches currently supported by IBM (the IBM 2042 family) are described in Table 4-5.

Table 4-5 The INRANGE switches

Model	Ports	Speed	Connectors	Rack space
2042-001 INRANGE FC/9000 Fibre Channel Director	64	1 Gb/s	SC	20U
2042-128 INRANGE FC/9000-128 Fibre Channel Director	128	1 Gb/s	SC	40U

In addition, IBM also sells the INRANGE FC/9000 40U cabinet (2042-C40). The cabinet is able to hold one 128-port switch or two 64-port switches.

Features and limitations

The INRANGE FC/9000 family of switches support both fabric-attached and loop-attached devices on any port.

Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on requirements.

For more information, see *IBM SAN Survival Guide Featuring the INRANGE Portfolio*, SG24-6150.

4.3.4 McDATA

The McDATA switches currently supported by IBM (the IBM 2031 and 2032 families) are described in Table 4-6.

Table 4-6 The McDATA switches

Model	Ports	Speed	Connectors	Rack space
2031-016 McDATA ES-3016 Fabric Switch	16	1 Gb/s	LC	1U
2031-032 McDATA ES-3032 Fabric Switch	32	1 Gb/s	LC	1.5U
2031-L00 McDATA ES-1000 Loop Switch	8+1	1 Gb/s	SC	1U
2032-001 McDATA ED-5000 Enterprise Fibre Channel Director	32	1 Gb/s	SC	18U
2032-064 McDATA ED-6064 Enterprise Fibre Channel Director	64	1 Gb/s	LC	9U

In addition, IBM also sells the McDATA FC-512 Cabinet (2032-C36).

Features and limitations

The only McDATA switch that supports loop attached devices (tape drives, for example) is the McDATA ES-1000 Loop Switch.

The ports in the 2032-001 and 2032-064 can be added in cards. The cards are not compatible between the models. Each card contains four ports. There are three different cards available:

- ▶ Four shortwave (multimode) ports
- ▶ Four longwave (singlemode) ports
- ▶ Three shortwave ports and one longwave port

On the 2031-016 and 2031-032 switches, any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC.

The 2031-L00 switch has two different types of ports:

- ▶ Eight switching Hub ports that can only be configured with shortwave (multimode) GBIC (or copper interface)
- ▶ One bridge port, for connecting the loop switch to the fabric, that can be configured with a shortwave (multimode) or longwave (singlemode) GBIC.

The McDATA switches support fabrics of up to 31 switches only.

For more information, see *IBM SAN Survival Guide Featuring the McDATA Portfolio*, SG24-6149.

4.4 Storage

The IBM Fibre Channel attached storage systems supported on HP 9000 series machines are described in Table 4-7 and Table 4-8.

Table 4-7 IBM Fibre Channel attached disk storage systems

Storage system	Notes
IBM TotalStorage Enterprise Storage Server	Only supports HP-UX 11.0
IBM TotalStorage FASiT500 Storage Server	
IBM TotalStorage FASiT200 Storage Server	Only supports IBM 2109 switch fabrics
IBM TotalStorage SAN Controller 160 with IBM 7133 disk subsystems	Only supports IBM 2109 switch fabrics

Table 4-8 IBM Fibre Channel attached tape storage systems

Storage system	Notes
IBM TotalStorage Enterprise Tape System 3590	Requires feature 3510 or 9510 on the drive Only supports IBM 2109 switch fabrics Cascaded switches are not supported
IBM 3584 UltraScalable Tape Library	Requires feature 1456 or 1466 on the drive Only supports IBM 2109 switch fabrics

Other storage system vendors also support their storage systems on HP 9000 series systems.

4.5 Installation procedures

In the topics that follow we detail the steps that need to be taken to install and successfully connect an HP-UX system.

4.5.1 Preparing to install the FC HBA

There are a number of checks that you need to ensure, at a minimum:

- ▶ Check that you have a free, supported PCI slot for the FC HBA.
- ▶ Check that you have the correct Processor Dependent Code (PDC) firmware installed, if you want to boot from a disk attached to the FC HBA.

If you need to upgrade the PDC version on your system, contact your HP representative.

- ▶ Check that you have the correct version and latest maintenance level of HP-UX installed, the minimum required levels are specified in Table 4-9.

Table 4-9 Required patches for HP-UX versions

HP-UX version	Required patches
10.20	J3630BA FCMS PHSS_23581 or later PHKL_17590 PHKL_16751
11.0	FCMS PHKL_21381
11i	FCMS PHKL_23626

- ▶ Obtain the latest drivers.
Download system firmware, HP-UX fixes, and maintenance levels from:
<http://www.software.hp.com/>

We use the following hardware and software in our installation:

- ▶ HP 9000 model L1000
- ▶ HP-UX 11.0
- ▶ HP A5158A Fibre Channel adapter
- ▶ Latest device driver updates for the Fibre Channel adapter

4.5.2 Installing the FC HBA

Our system has HP-UX preloaded at the factory and the Fibre Channel adapter is already installed. Because of this, we skip steps 1-6.

1. Power on the system, if it is not already on, and log in as user root.
2. Check that you have the correct HP-UX maintenance level installed.
3. Install the base level of device drivers for your FC HBA.
4. Update the drivers you just installed to the latest versions.
5. Shut down the system.
6. Install the host adapter card in a supported slot.

Depending on the model of your system, the adapter is installed by HP's Service Support Representative (SSR), or you need to install it yourself. You can find the installation instructions in the documentation provided with your system.

Note: Some HP 9000 models support hot-pluggable adapters. If your system is one of those models, you have the option of installing the adapter without shutting down the system. If you want to make use of the hot plug feature, follow the directions in the installation manual delivered with your system.

7. Connect the adapter to your Fibre Channel fabric.
8. Power on the system and log in as user root.
9. Check that the adapter is correctly configured as in Example 4-1.

Example 4-1 Checking the FC HBA

```
# ioscan -funC fc
Class      I H/W Path  Driver S/W State  H/W Type      Description
=====
fc         0 0/4/0/0  td  CLAIMED  INTERFACE    HP Tachyon TL/TS Fibre
Channel Mass Storage Adapter
                        /dev/td0
#
```

You get a list of the Fibre Channel cards installed in the system.

The adapter is correctly configured by the operating system, if the S/W State is CLAIMED.

10. Check the parameters of your adapter as in Example 4-2

Example 4-2 Checking the FC HBA parameters

```
# fcmsutil /dev/td0

Vendor ID is = 0x00103c
Device ID is = 0x001028
TL Chip Revision No is = 2.3
PCI Sub-system Vendor ID is = 0x00103c
PCI Sub-system ID is = 0x000006
Topology = PTTOPT_FABRIC
Local N_Port_id is = 0x021700
N_Port Node World Wide Name = 0x50060b00000781df
N_Port Port World Wide Name = 0x50060b00000781de
Driver state = ONLINE
Hardware Path is = 0/4/0/0
Number of Assisted IOs = 1188
Number of Active Login Sessions = 0

#
```

You can find the World Wide Name (WWN) of the adapter as the value of the N_Port Port World Wide Name parameter. You need to know the WWN to configure your storage system.

You can find the protocol the adapter is using as the Topology parameter. It should be PTTOPT_FABRIC for switched fabric attachment.

You can also see, where in the switch fabric our card is attached to, by looking at the hexadecimal Local N_Port_id is parameter. In our case it is 0x021700. From the leftmost byte, byte 1 (02) is the domain id of the switch this adapter is connected to, the third byte is '17' where '1' is the switch board within the switch, and '7' is the port number within the switch board. Byte 3 is always '00' for a switch-attached adapter.

11. Optionally, check to see that the port is correctly configured on the switch

The procedure you need to use is different for each switch vendor. We show the procedure for the IBM 2109 family of switches in Example 4-3. For INRANGE switches, you should use the IN-VSN Enterprise Manager application to find the same information. For McDATA switches, you can find the information in the EFC Manager application.

Example 4-3 Checking the FC HBA connectivity to the switch

```
# telnet 9.1.38.58
Trying...
Connected to 9.1.38.158.
Escape character is '^]'.

Fabric OS (tm) Release a2.4.1c

Login: admin
Password:

itsosw2:admin> switchShow
switchName: itsosw2
switchType: 3.4
switchState: Online
switchRole: Subordinate
switchDomain: 2
switchId: fffc02
switchWwn: 10:00:00:60:69:20:1d:74
switchBeacon: OFF
port 0: id Online      E-Port 10:00:00:60:69:20:1d:4e "itsosw1" (upstream)
port 1: id No_Light
port 2: -- No_Module
port 3: id No_Light
port 4: id Online      F-Port 50:00:1f:e1:00:0a:e4:02
port 5: id Online      F-Port 10:00:00:00:c9:20:d2:1f
port 6: -- No_Module
port 7: id Online      F-Port 50:06:0b:00:00:07:81:de
itsosw2:admin> quit
Connection closed.
#
```

From Example 4-3, we see that the port we are using (port 7) is online, it is an F_Port (for fabric attachment), and the WWN listed is the same that we see on our adapter in Example 4-2.

12. Your adapter is now correctly installed.

4.5.3 Installing the storage system

After the adapter installation, we have to create any necessary switch zones and allocate the storage for our host from the storage system we are using. As both of those steps are different for different switch manufacturers and different storage subsystems, they are not described here.

Note: The procedure described below is very generic. You may have to install additional device drivers depending on your storage system, like the Host Attachment Scripts for ESS or the atdd driver for IBM tape drives. Refer to the documentation that came with your storage system.

We have access to a single volume in a storage subsystem. We want to configure that volume to our host system. The procedure is as follows:

1. Scan for new devices:

```
# ioscan -f >/dev/null
#
```

Note: The `ioscan -f` command lists all the devices in the system, producing a very long listing. We are only interested in certain devices, and we check for them later.

2. Check that the switches (switch domains) are found:

```
# ioscan -funC fcp
Class      I  H/W Path  Driver S/W State  H/W Type  Description
-----
fcp        2  0/4/0/0.1  fcp  CLAIMED      INTERFACE  FCP Domain
fcp        1  0/4/0/0.2  fcp  CLAIMED      INTERFACE  FCP Domain
#
```

3. Check that the new storage device (disk in our case) is found:

```
# ioscan -funC disk
Class      I  H/W Path  Driver S/W State  H/W Type  Description
-----
disk       0  0/0/1/1.2.0  sdisk CLAIMED      DEVICE    SEAGATE ST39204LC
           /dev/dsk/c1t2d0 /dev/rdisk/c1t2d0
disk       1  0/0/2/0.2.0  sdisk CLAIMED      DEVICE    SEAGATE ST39204LC
           /dev/dsk/c2t2d0 /dev/rdisk/c2t2d0
disk       2  0/0/2/1.2.0  sdisk CLAIMED      DEVICE    HP        DVD-ROM 304
           /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0
disk       3  0/4/0/0.1.17.255.0.0.0  sdisk CLAIMED      DEVICE    DEC      HSG80
#
```

4. If the new disk is missing the special files (`/dev/dsk/*` and `/dev/rdisk/*`) like above, create them as follows:

```
# mksf -C disk -H 0/4/0/0.1.17.255.0.0.0
# mksf -C disk -H 0/4/0/0.1.17.255.0.0.0 -r
#
```

You need both the block device and the character device.

Check the storage devices again, as follows:

```

# ioscan -funC disk
Class I H/W Path Driver S/W State H/W Type Description
-----
disk 0 0/0/1/1.2.0 sdisk CLAIMED DEVICE SEAGATE ST39204LC
/dev/dsk/c1t2d0 /dev/rdisk/c1t2d0
disk 1 0/0/2/0.2.0 sdisk CLAIMED DEVICE SEAGATE ST39204LC
/dev/dsk/c2t2d0 /dev/rdisk/c2t2d0
disk 2 0/0/2/1.2.0 sdisk CLAIMED DEVICE HP DVD-ROM 304
/dev/dsk/c3t2d0 /dev/rdisk/c3t2d0
disk 3 0/4/0/0.1.17.255.0.0.0 sdisk CLAIMED DEVICE DEC HSG80
/dev/dsk/c12t0d0 /dev/rdisk/c12t0d0
#

```

5. Start using the device like any disk:

We use the procedure shown below to create a volume group, small logical volume, and filesystem on the new disk and test it, as in Example 4-4.

Example 4-4 Creating volume group and filesystem on the new disk

```

# mkdir /dev/vg01
# mknod /dev/vg01/group c 64 0x010000
# pvcreate /dev/rdisk/c12t0d0
Physical volume "/dev/rdisk/c12t0d0" has been successfully created.
# vgcreate -s 16 /dev/vg01 /dev/dsk/c12t0d0
Volume group "/dev/vg01" has been successfully created.
Volume Group configuration for /dev/vg01 has been saved in /etc/lvmconf/vg01.conf
# lvcreate -L 32 /dev/vg01
Logical volume "/dev/vg01/lvol1" has been successfully created with
character device "/dev/vg01/rlvol1".
Logical volume "/dev/vg01/lvol1" has been successfully extended.
Volume Group configuration for /dev/vg01 has been saved in /etc/lvmconf/vg01.conf
# newfs /dev/vg01/rlvol1
newfs: /etc/default/fs is used for determining the file system type
version 3 layout
32768 sectors, 32768 blocks of size 1024, log size 1024 blocks
unlimited inodes, 32768 data blocks, 31672 free data blocks
1 allocation units of 32768 blocks, 32768 data blocks
first allocation unit starts at block 0
overhead per allocation unit is 0 blocks
# mkdir /testfilesystem
# mount /dev/vg01/lvol1 /testfilesystem
# df -k /testfilesystem
/testfilesystem (/dev/vg01/lvol1) : 30797 total allocated Kb
29688 free allocated Kb
1109 used allocated Kb
3 % allocation used
# cp /etc/hosts /testfilesystem
# ls -l /etc/hosts /testfilesystem/hosts
-rw-r--r-- 1 root sys 578 Sep 26 15:27 /etc/hosts
-rw-r--r-- 1 root sys 578 Nov 28 16:44 /testfilesystem/hosts
# diff /etc/hosts /testfilesystem/hosts
#

```

The connection is successful.



NUMA-Q Fibre Channel implementation

In this chapter we discuss how to successfully connect IBM NUMA-Q platforms running Dynix/ptx to the SAN fabric. We cover the step-by-step procedure from start to finish, and the supported products which will work with the NUMA-Q system.

The following topics are discussed:

- ▶ “Overview of the NUMA-Q system” on page 108
- ▶ “Fibre Channel host bus adapters” on page 109
- ▶ “Fibre Channel switch support” on page 112
- ▶ “Supported storage on the NUMA-Q fabric” on page 114
- ▶ “Installation planning” on page 116
- ▶ “Installation procedure” on page 117
- ▶ “Device functionality checks” on page 125
- ▶ “Specific limitations” on page 126

5.1 Overview of the NUMA-Q system

The NUMA-Q system is composed of four processor building blocks called “quads”. Each quad has two separate PCI busses for installing IO cards. You can interconnect the quads to each other to create a single larger system. When you connect the quads together in this manner, they run on a single instance of the operating system, and share their memory, storage, and networking capabilities.

The NUMA-Q system runs with only the Dynix/ptx operating system. A single instance of the operating system is shared across all of the quads configured in a single system node. There are three levels of Dynix/ptx that can be installed on the NUMA-Q system. These are Dynix/ptx V4.4.10, Dynix/ptx V4.5.3, or Dynix/ptx V4.6.1. Fibre Channel operations are generally the same on all of these operating systems, with only some differences in the level of feature support in some areas. These limitation differences will be discussed in their respective sections throughout this chapter.

The NUMA-Q system supports Fibre Channel Point-to-Point, Fibre Channel Arbitrated Loop (FC-AL), or Fibre Channel Switch (Switched) topologies. Only one topology is supported at any one time across the system.

The topology is selectable via the kernel parameter **ff_topology**, found in the file: **/usr/conf/uts/io/ff/ff_space.c**. This parameter must be set for the topology you desire to run. For point-to-point fabric it is recommended that you create a direct FC HBA to storage simple FC-AL with only two connections rather than using the point to point topology setting. NUMA-Q supports fabric levels 1, 2, and 3. The NUMA-Q system supports class 2, and class 3 Fibre Channel protocol. You can also establish a cascaded, and zoned fabric environment with the NUMA-Q.

The NUMA-Q system does not use the IBM SDD driver product, as the Dynix/ptx operating systems come with a built-in multipath driver. This driver will automatically try to use all paths it detects to the fibre devices. It also requires that all paths on a common fabric be able to see all of the system’s devices equally which are on the fabric. For this reason you must take special care when configuring devices on a fabric.

For additional documentation which would be of great value to you when performing SAN planning, and installation on the NUMA-Q; and product specific implementation documentation for supported direct fibre attached storage devices, see the following Web site:

<http://webdocs.numaq.ibm.com/>

From this location, select Fibre Channel, SAN, Disk Systems, or Tape Systems as appropriate.

5.2 Fibre Channel host bus adapters

With all the Dynix/ptx operating systems, there is driver support incorporated into the base operating system.

In Table 5-1 we list all of the Fibre Channel host bus adapters (FC HBA) which are supported by the NUMA-Q system.

Table 5-1 Supported FC HBAs

FC HBA type	Feature Code	IBM Part #	Usage comments
Emulex LP7000E (nicknamed Superfly)	7556	9P8178	This adapter supports all of the NUMA-Q supported storage subsystems, and quad types. It does not support any FC-AL hub connections.
Emulex LP6000 V3 non-OFC (nicknamed Firefly)	7599	9P8183	This is the most recent model of this adapter, and is used with NUMA-Q quads Q06 through Q11, and Mquad Q01 through Q05 based systems. Some of the newer storage subsystems are not supported with this adapter. It does not support any FC-AL hub connections.
Emulex LP6000	7598		Used only in quad based systems. Did not support switched environments. Supported FC-AL hubs. Some of the newer storage subsystems are not supported with this adapter.
Emulex LP6000	7597		Used only with NUMA-Q quads Q06 through Q11, Some of the newer storage subsystems are not supported with this adapter.

All of these FC HBAs are shipped with a multimode, shortwave GLM (Gigabit Link Module) installed. Some of these adapters work only in certain types of systems and environments. See “Fibre Channel/x430/NUMA-Q Service/Host Adapter Overview” for further details at the Web site:

<http://webdocs.numaq.ibm.com/>

When installing the FC HBA(s), it is important that you first know how many you will be installing. With multiple quad systems it is strongly recommended that you install one FC HBA per fabric in each of the quads making up the system; this is for performance and redundant paths.

Therefore for a four quad system with a single fabric, you should install four adapters; with two fabrics you would install eight; and so on.

In Figure 5-1 we show the LP7000E adapter, and in Figure 5-2 we show the LP 6000 adapter.

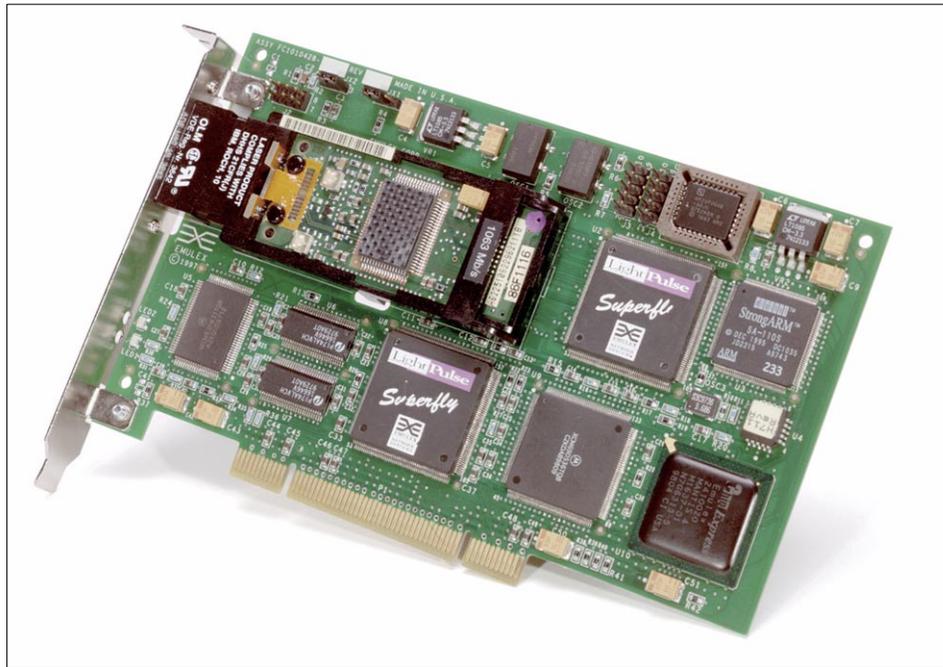


Figure 5-1 Emulex LP7000E Fibre Channel Adapter

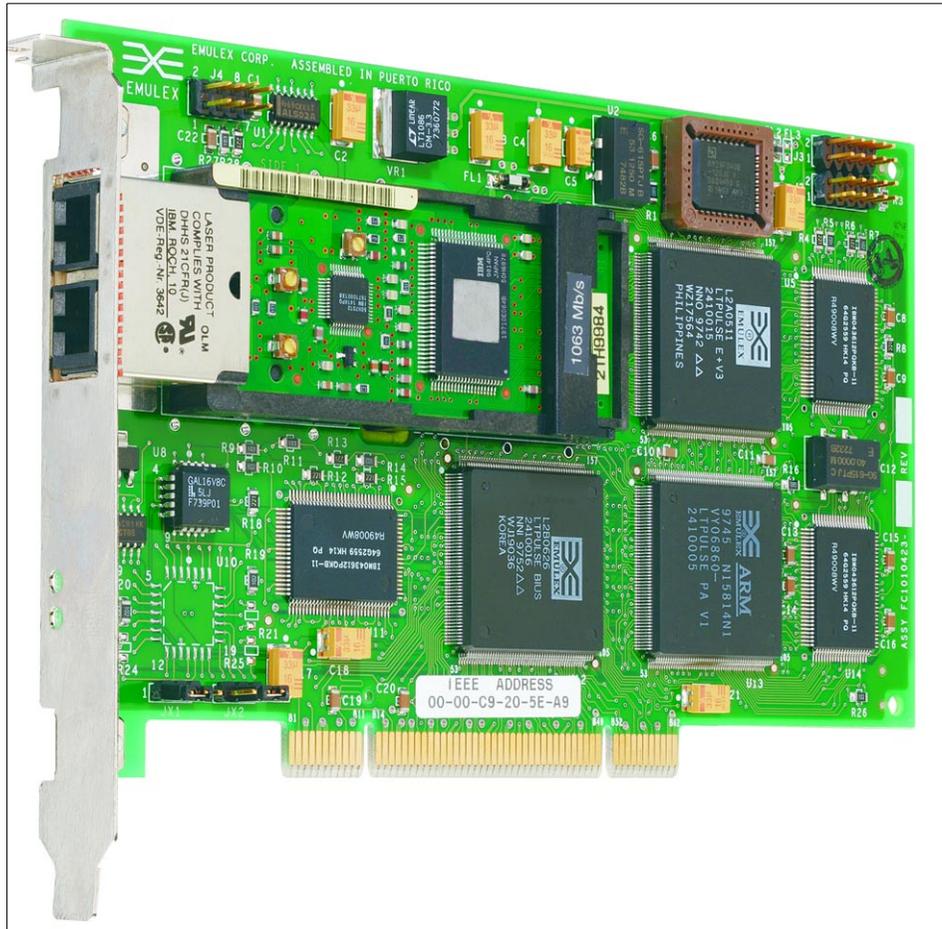


Figure 5-2 Emulex LP6000 Fibre Channel Adapter

5.2.1 Host adapter placement

The first FC HBA installed in a NUMA-Q system should be installed in PCI Bus 1 slot 6, of the boot quad (usually quad 0). If a second adapter is to be added to the same quad it should be installed in PCI Bus 0 slot 3. Further FC HBAs being added should be distributed evenly between the two PCI busses. A maximum of four FC HBAs can be installed in a single quad.

You must not install two FC HBAs to the same fabric in the same quad.

5.2.2 Host adapter installation preparation

The firmware to use for the FC HBA is provided on the Dynix/ptx Base and Layered Products CD, which is provided with the NUMA-Q system. It is loaded onto the NUMA-Q system using the standard layer product installation model defined in the appropriate documentation for the level of operating system you are running. The only supported version of the firmware is that provided by NUMA-Q, as the driver is only proofed with those versions. For this reason we recommend that you use these versions only.

5.3 Fibre Channel switch support

The NUMA-Q system is tested and proofed with the switches outlined in Table 5-2.

Table 5-2 The NUMA-Q supported switch family

Model	Ports	Speed	Connectors	Rack space
Brocade Silkworm 1000 switch	6 or 16	1 Gb/s	SC	2U
Brocade Silkworm 2000 switch	8 or 16	1Gb/s	SC	2U
IBM TotalStorage SAN Switch S08	8	1 Gb/s	SC	1U
IBM TotalStorage SAN Switch S16	16	1 Gb/s	SC	2U

These are the only switches which are formally supported with the NUMA-Q system. For host attachment the switch must use the multimode shortwave laser GBICs (black in color). This allows for distances up to 500 meters. For long distances up to 10Km, two switches can be used with singlemode longwave GBICs (light blue in color) connecting them together.

In Figure 5-3 we show the IBM 2109-S08 and S16 switches.



Figure 5-3 IBM 2109-S08 (top) and S16

5.3.1 Fibre Channel switch installation preparation

If you are installing the NUMA-Q on a new, unused Fibre Channel switch you should set up the fabric as a single zone, and familiarize yourself with a simple network before expanding it. If, however, you are adding the NUMA-Q to an already existing fabric, you may need to configure the switch to work with multiple hosts, different vendors and operating systems. For this more planning is required.

Using zones with the switches

When you are sharing a switch with other hosts systems in a heterogeneous environment, the NUMA-Q requires its own zone to be established. The NUMA-Q cannot share a zone with any other hosts in a heterogeneous fabric. It can share the switch as long as each host is zoned separately. It is strongly recommended that these zones be established using hardware (port number) zoning for the host ports to provide increased security for the host attachments. The process for establishing zones is covered in detail in SAN/Zoning/Concepts at the following Web site.

<http://webdocs.numaq.ibm.com/>

Using cascaded switches

The NUMA-Q system can also be connected to a cascaded, or meshed fabric. Some limitations exist with certain releases of the operating system in this area. See Section 5.8, “Specific limitations” on page 126 to ensure your planned environment will work with the version of Dynix/ptx you are using. Further details on establishing cascaded fabrics is covered in detail in the multi-switch sections of SAN/x430/NUMA-Q/Configurations at:

<http://webdocs.numaq.ibm.com/>

5.4 Supported storage on the NUMA-Q fabric

The NUMA-Q system supports a great number of storage devices on its fabric. It also has its own fibre to SCSI bridge (IBM feature code 3019) allowing it to connect a number of SCSI devices as well.

For some SCSI tape support the SAN Data Gateway (SDG) IBM 2108 is used. Table 5-3 is a list of the standard supported direct fibre disk devices that are available for NUMA-Q:

Table 5-3 IBM Fibre Channel attached disk storage systems

Storage system	Method of attachment	Operating systems
IBM 3542/3560 - FAST200 Storage Server with up to two EXP500 disk expansion unit	Attach by both fibre switch, or FC-AL.	Supported with DYNIX/ptx V4.4.8 and higher, V4.5.2 and higher, and V4.6.1.
IBM 2105 Enterprise Storage Server (ESS)	Switched or FC-AL direct attached from the FC HBA; maximum of 128 LUNs per LSS supported with DYNIX/ptx V4.4.10	Supported with DYNIX/ptx V4.5.2 and higher, and V4.6.1, V4.4.10 (see note below).
EMC Symmetrix Storage Subsystems - Switched or FC-AL direct attached from the FC HBA	Switched or FC-AL direct attached from the FC HBA	Supported with DYNIX/ptx V4.5.2 and higher, and V4.6.1, V4.4.4 and higher.
EMC/CLARiiON DAE (Direct Attach Enclosure) JBOD (Just a Bunch Of Disk)	FC-AL direct attached from the FC HBA only	Supported with DYNIX/ptx V4.5.1 and higher, and V4.6.1, V4.4.7 and higher.

In addition to the above standard supported list of disk devices, there is also a list of Request for Price Quotation (RPQ) disk devices that are available. These are listed in Table 5-4.

Table 5-4 IBM NUMA-Q RPQ Fibre Channel attached disk storage systems

Storage system	Method of attachment	Operating systems
IBM 2105 Enterprise Storage Server (ESS)	Switched or FC-AL direct attached from the FC HBA.	DYNIX/ptx V4.4.7/8/9.
HDS Freedom 9960	Switched or FC-AL direct attached from the FC HBA	Supported with DYNIX/ptx V4.5.2 and higher, and V4.6.1, V4.4.7 and higher.
HDS Freedom 9200 Series	Switched or FC-AL direct attached from the FC HBA	Supported with DYNIX/ptx V4.5.2 and higher, and V4.6.1, V4.4.7 and higher.
HDS 7700E	Switched or FC-AL direct attached from the FC HBA	Supported with DYNIX/ptx V4.5.1 and higher, and V4.6.1, V4.4.4 and higher.
HDS 5800 Series	Switched or FC-AL direct attached from the FC HBA	Supported with DYNIX/ptx V4.5.1 and higher, and V4.6.1, V4.4.4 and higher.
HDS 5700F	Switched or FC-AL direct attached from the FC HBA	Supported with DYNIX/ptx V4.5.1 and higher, and V4.6.1, V4.4.4 and higher.

Table 5-5 describes the standard supported direct fibre tape device that is available for NUMA-Q, with details.

Table 5-5 IBM NUMA-Q standard Fibre Channel attached tape storage systems

Storage system	Method of attachment	Operating systems
IBM 3590E tape drives with direct fibre interface	FC-AL attached from either the FC HBA directly, or a FL port on a Brocade switch	Supported with DYNIX/ptx V4.5.2 and greater.

In addition to the above standard supported tape device, there is also a table of the Request for Price Quotation (RPQ) tape device that is available. See Table 5-6 for details.

Table 5-6 IBM NUMA-Q RPQ Fibre Channel attached tape storage systems

Storage system	Method of attachment	Operating systems
STK 9840 tape drives with direct fibre interface	FC-AL attached from either the FC HBA directly, or a FL port on a Brocade switch	Supported with DYNIX/ptx V4.5.2 and greater.

5.5 Installation planning

Before starting the installation, it is important to plan what you intend to do. This includes deciding what will be connected and where they will be located. Our example install will be small, although your installation may need to be quite a lot larger.

Items to consider will be the physical location of your system, the storage that is being connected, and the switch or switches which you are connecting them to. Cables must be purchased to connect these all together; or possibly an installer will be called in if running between floors of a building, or over long distances. Ready made cable lengths are available for order for most general installations. Table 5-7 shows the cable lengths and ordering information for these cables.

Table 5-7 Ready made fibre optic duplex cables

Length	Multimode style model number	Singlemode style model number
3 meters	CBL-0300-F3	CBL-0300-F2
8 meters	CBL-0800-F3	CBL-0800-F2
20 meters	CBL-2000-F3	CBL-2000-F2
35 meters	CBL-3500-F3	CBL-3500-F2
70 meters	CBL-7000-F3	CBL-7000-F2

With fibre cables, remember that attenuation is your enemy; it can cause your entire network to be highly unstable, or not to function at all. Watching your number of connections, ensuring clean ends, and proper seating, is critical to your success. Do not try to stretch the maximum distances of the signal.

See Table 5-8 for types and maximum limits.

Table 5-8 Common Fibre Channel cabling configurations

GBIC type	Cable type	Distance (1 Gb/s)	Distance (2 Gb/s)
longwave	9 micron, singlemode	10 km	10 km
shortwave	50 micron, multimode	500 m	300 m
shortwave	62.5 micron, multimode	300 m	150 m

Ensure that any cables you buy or have made are constructed with the correct connector ends to match your switch ports and FC HBAs. All NUMA-Q 1Gb/s cabling uses the ISC type connectors.

Note: NUMA-Q does not support the 2 Gb/s Fibre Channel speed.

5.6 Installation procedure

In this section we run through a simple fabric installation and show the success of the work. We are using a two quad NUMA-Q system with four FC HBAs installed, connected to an IBM ESS using two FC host adapters. We have four LUNs configured on the ESS for access to the NUMA-Q system. We are also using Dynix/ptx V4.5.3 to perform this installation. Any special instructions which may differ from our performed steps for the other supported releases of the operating system are noted as we progress.

To start with, it is assumed that this is not a new installation, and that the Fibre Channel we are installing is an add-on to the platform hardware. The system is powered up and running the operating system of choice. It is also assumed that the system has been configured for LAN and TCP/IP support, and has the supported firmware package installed. Now we are ready to start installing our fabric.

5.6.1 Verifying the correct software is installed

1. Before we start our install, we make one last check to ensure the proper software is really on the system. This we do by running:

```
# cat /etc/versionlog
Sun Jan 14 17:42:33 PST 2001 - DYNIX/ptx(R) Base V4.5.3 PN:
1003-77788-03_A00
Sun Jan 14 17:42:45 PST 2001 - DYNIX/ptx(R) Encryption V4.5.3 PN:
1003-77790-03_A00
```

```

Sun Jan 14 17:45:51 PST 2001 - DYNIX/ptx(R) Public Software V4.5.3 PN:
1003-77789-03_A00
Sun Jan 14 17:45:58 PST 2001 - ptx/INSTALL V2.0.4 PN: 1003-78071-04_A00
Sun Jan 14 17:46:32 PST 2001 - ptx/BaseComms V1.2.0 PN: 1003-77812-00_A00
Sun Jan 14 17:47:48 PST 2001 - ptx/LAN V4.7.1 PN: 1003-77834-01_A00
Sun Jan 14 17:50:57 PST 2001 - ptx/TCP/IP V4.6.2 PN: 1003-77894-02_A00
Sun Jan 14 17:53:00 PST 2001 - ptx/EFS V1.4.2 PN: 1003-77828-02_101
Sun Jan 14 17:53:02 PST 2001 - Layered product Install Patch V1.0.2 PN:
1003-78507-02_A00
Sun Jan 14 17:54:08 PST 2001 - Online Diagnostic Software V2.7.2 PN:
1003-76392-02_A00
Sun Jan 14 17:54:27 PST 2001 - ptx/EES V1.2.2 PN: 1003-78904-02_A00
Sun Jan 14 18:05:59 PST 2001 - ptx/XWM V4.6.2 PN: 1003-77900-02_A00
Sun Jan 14 18:15:56 PST 2001 - ptx/JSE V3.0.2 PN: 1003-79941-00_A00
Sun Jan 14 18:16:28 PST 2001 - ptx/edb V3.4.2 PN: 1003-79398-02_A02
Sun Jan 14 18:18:33 PST 2001 - ptx/NFS V4.7.1 PN: 1003-77841-01_A00
Sun Jan 14 18:19:50 PST 2001 - ptx/Configuration Assistant V1.0.0 PN:
1003-79203-00_A02
Sun Jan 14 18:20:37 PST 2001 - ptx/SPDRIVERS V3.2.0 PN: 1003-81493-00_A00
Sun Jan 14 18:24:29 PST 2001 - CFWDL FW Bundle V1.2.0 PN: 1003-81485-00_A00
Mon Jan 22 15:54:10 PST 2001 - ptx/JSE v1.1.2 PN: 1003-74173-03_A00

```

From this list we can see that the operating system is at a supported level, and that the necessary LAN products, and CFWDL FW Bundle are installed as well.

2. As our system was powered up and running, we now shut it down so we can install the add-on FC HBAs into the system.

5.6.2 FC HBA setup

3. We install the FC HBA into the appropriate quad0 slot6 of PCI bus 1 for our system configuration. Remember that this will vary with larger systems.
4. After installing the FC HBA we repower the system and boot it up. At this time the boot probe has discovered our new FC HBA and we can confirm it is present and gather data on it by using the following commands.

```

$ /etc/dumpconf | grep ff
quad0      quad      0 0xff000000 L   sci      sci0
asy0       asy       0 0xff000000 L   eisa     quad0
asy1       asy       1 0xff000001 L   eisa     quad0
mdc0       mdc       0 0xff000000 L   eisa     quad0
ff0        ff        0 0x00000000 SP  pci      quad0
sd0        sd        0 0xff000000 L   scsi     scsibus0
ff1        ff        1 0x00000006 SP  pci      quad0
quad1      quad      1 0xff000001 L   sci      sci0
+ff2       ff        2 0x00000000 SP  pci      quad1
fabric3    fabric    3 0x00011c00 SMX  fc       ff2
ff3        ff        3 0x00000006 SP  pci      quad1

```

5. Now we run **devctl -N** command to assign the temporary devices to their permanent names.

```
# devctl -N
devctl: Assigned +ff2 --> ff2
```

6. Next we gather data from it using the **infodev -a** command. This command provides us with the WWN which we will need when configuring the LUNs on the ESS to be assigned to the NUMA-Q system.

```
$ /etc/infodev -a ff2
Vendor ID                : Emulex
Product ID               : Firefly
Revision Level           : SuperFly-4
Device Class              : ctlr
Node WWN                 : 10:00:00:00:c9:21:58:59
Port WWN                  : 10:00:00:00:c9:21:58:59
Located on                : quad1
```

7. Finally, we check to ensure the correct version of FC HBA firmware is loaded on the adapter. If this is not the correct version we would use the same **cfwdl** command with the download option to load the correct firmware file. See the **cfwdl** manpage for details on the command usage.

```
# /usr/service/bin/cfwdl -a -d ff2

cfwdl: Versions for device /dev/ff/ff2:
Chipset Revision Levels:
    BIU Revision = 0x4
    Sequence Manager Revision = 0x10020193
    Endec Revision = 0x3001506d

Firmware Revisions Levels:
    Operational Firmware Revision = 0x2232000(2.23 0)
    Diagnostic Firmware Revision = 0x1212000(1.21 0)

FC-PH Version Supported:
    Highest FC-PH Version Supported = 4.3
    Lowest FC-PH Version Supported = 4.3

cfwdl: Flash program list versions for: /dev/ff/ff2
    Diagnostic (boot) firmware version: 1.21
    Test firmware version: 1.20
    Functional firmware version: 2.23
    SLI-1 firmware version: 2.23
    SLI-2 firmware version: 2.23
Device Type is SuperFly/LP7000
```

5.6.3 Switch setup

8. Using the switch's *local control panel*, we set the switch to its default settings and password; and also set the TCP/IP address as needed to be able to access it using the remote telnet login to complete the configuration as needed. This IP address can also be used to access the switch via the Web based GUI session.
9. Now we telnet to the switch with the IP address and login using the default login, and password (login=admin, password= password).

```
# telnet 9.43.227.19
Trying...
Connected to 9.43.227.19.
Escape character is '^]'.
```

```
Fabric OS (tm) Release a2.4.1c
```

```
login: admin
Password:
```

```
SNJ2109DD2:admin>
```

10. We also set the switch for a unique name, correct date and time, and change the following two settings as required for the NUMA-Q:
 - The fabric parameter `route.delayReroute` must be disabled.
 - The fabric parameter `route.stickyRoutes` must be enabled.

If any additional help is needed with these settings refer to Fibre Channel/IBM 2109 or SilkWorm 2000/Configuration at url:

<http://webdocs.numaq.ibm.com/>

11. We confirmed these settings by running the following **configShow** command.

```
SNJ2109DD2:admin> configShow
Ethernet addresses: 0:60:69:12:3f:fc
Nvram datadata: fei(0,0)host:/usr/switch/firmware e=9.43.227.19:fffffc00
g=9.43.
224.1 u=user tn=SNJ2109DD2
diag.postDisable:      0
fabric.domain:         1
fabric.ops.BBCredit:   16
fabric.ops.E_D_TOV:    2000
fabric.ops.R_A_TOV:    10000
fabric.ops.dataFieldSize: 2112
fabric.ops.mode.SeqSwitching: 0
fabric.ops.mode.fcpProbeDisable: 0
fabric.ops.mode.isolate: 0
fabric.ops.mode.longDistance: 0
```

```

fabric.ops.mode.noClassF:      0
fabric.ops.mode.pidFormat:    0
fabric.ops.mode.sync:         0
fabric.ops.mode.tachyonCompat: 0
fabric.ops.mode.unicastOnly:  0
fabric.ops.mode.useCsCtl:     0
fabric.ops.mode.vcEncode:     0
fabric.ops.vc.class.2:        2
fabric.ops.vc.class.3:        3
fabric.ops.vc.config:         0xc0
fabric.ops.vc.linkCtrl:       0
fabric.ops.vc.multicast:      7
fc4.fcIp.address:             0.0.0.0
fc4.fcIp.mask:                0.0.0.0
fcAL.fanFrameDisable:        0
fcAL.openSendCLS:            0
fcAL.useAltBBCredit:         0
lcdContrast:                  128
lcdContrast.orange:          208
licenseKey:                   zdeSSR9zy2efTTY;SbRcyRQScdTcSp0y
route.delayReroute:           0
route.embeddedPortBcast:      1
route.stickyRoutes:           1
rpc.rapid:                    1
rpc.rstatd:                   0
rpc.rusersd:                  0
shell.delete:                 0
shell.quiet:                  0
shell.timeout:                0
syslog.errDisplayLevel:       5
syslog.errsOnConsole:         1
thresh.alarmsfilter:          1
thresh.thad:                  1
xrelativeModeDisable:         0
Displaying enabled/disabled portlog events
Type <CR> to continue, Q<CR> to stop: q

```

12. We now need to check to ensure we are loaded with the proper switch firmware. This we do by running the **version** command on the switch telnet session.

```

SNJ2109DD2:admin> version
Kernel:      5.3.1
Fabric OS:   a2.2.1a
Made on:     Tue May 8 15:44:06 PDT 2001
Flash:      Tue May 8 15:45:24 PDT 2001
BootProm:   Thu Jun 17 15:20:39 PDT 1999

```

13. We now need to connect the cables ensuring that we are careful not to touch the ends and making sure they are clean of any dirt or dust. We seat them securely into the GLM of the FC HBA, and the GBIC of the switch port. By issuing the **switchShow** command we can see these switch details:

```
SNJ2109DD2:admin> switchShow
switchName:    SNJ2109DD2
switchType:    2.4
switchState:   Online
switchRole:    Principal
switchDomain:   1
switchId:      fffc01
switchWwn:     10:00:00:60:69:12:3f:fc
switchBeacon:  OFF
port 0: sw No_Light
port 1: sw No_Light
port 2: id No_Light
port 3: sw No_Light
port 4: -- No_Module
port 5: -- No_Module
port 6: -- No_Module
port 7: -- No_Module
port 8: -- No_Module
port 9: -- No_Module
port 10: -- No_Module
port 11: -- No_Module
port 12: id Online      F-Port 10:00:00:00:c9:21:58:59
port 13: sw No_Light
port 14: sw No_Light
port 15: id No_Light
```

5.6.4 Storage setup

14. Now we need to ensure our storage device is ready to be connected to our NUMA-Q system. In our example we are using a group of LUNs from an IBM ESS. So, we must now define the FC HBA to the ESS using the StorWatch manager. As we are sharing the ESS between multiple hosts and the NUMA-Q, the ESS ports will be configured for Access Restricted, and specific LUN assignments will be used. To do this we must map the desired LUNs to the FC HBAs WWN. We get the Node WWN of the FC HBA from the **infodev** command we ran in step 3 above.
15. We now need to connect the cables from the ESS to the switch, again being careful not to contaminate them with dirt or oil from our hands. We seat them securely into the GLM of the ESS adapter, and the GBIC of the switch port. Another check using **switchShow** gives us the following:

```
SNJ2109DD2:admin> switchShow
```

```

switchName:    SNJ2109DD2
switchType:    2.4
switchState:   Online
switchRole:    Principal
switchDomain:   1
switchId:      fffc01
switchWwn:     10:00:00:60:69:12:3f:fc
switchBeacon:  OFF
port 0: sw No_Light
port 1: sw No_Light
port 2: id No_Light
port 3: sw No_Light
port 4: -- No_Module
port 5: -- No_Module
port 6: -- No_Module
port 7: -- No_Module
port 8: -- No_Module
port 9: -- No_Module
port 10: sw Online      F-Port 10:00:00:00:c9:22:a2:8d
port 11: sw Online      F-Port 10:00:00:00:c9:28:05:da
port 12: id Online      F-Port 10:00:00:00:c9:21:58:59
port 13: sw No_Light
port 14: sw No_Light
port 15: id No_Light

```

16. After connecting the cables we are ready to probe for the LUNs. This is done by using the **devctl -c ff2** command on the NUMA-Q system, which we will point at the FC HBA (ff2).

```

# devctl -c ff2
devctl: Found +fabric3, +sd713, +sd714, +sd715, +sd716,

```

This probe should find the fabric, and all of the LUN devices that we assigned to the FC HBA WWN on the ESS. We see that these devices are all assigned names with a “+” sign attached to them. This is a temporary assigned name, which we now need to make permanent.

17. Now we again run **devctl -N** command to assign the temporary devices to their permanent names.

```

# devctl -N
devctl: Assigned +sd713 --> sd713
devctl: Assigned +sd714 --> sd714
devctl: Assigned +sd715 --> sd715
devctl: Assigned +sd716 --> sd716

```

18. We run **/etc/dumpconf -d** command to see all the devices now found on the system and their parent tree.

```

# dumpconf -d
NAME          CFGTYPE  DEVNUM  UNIT          FLAGS  OnBUS  OnDEVICE
quad0         quad     0       0xff000000   L      sci    sci0

```

asy0	asy	0	0xff000000	L	eisa	quad0
asy1	asy	1	0xff000001	L	eisa	quad0
mdc0	mdc	0	0xff000000	L	eisa	quad0
ff0	ff	0	0x00000000	SP	pci	quad0
qlc0	qlc	0	0x00000001	L	pci	quad0
scsibus0	scsibus	0	0x00000070	L	m SCSI	qlc0
sd0	sd	0	0xff000000	L	scsi	scsibus0
tm0	tm	0	0x00000040	L	scsi	scsibus0
cd0	cd	0	0x00000050	L	scsi	scsibus0
pci0	pci	0	0x00000003	L	pci	quad0
pe0	pe	0	0x00000003	L	pci	pci0
pe1	pe	1	0x00000003	L	pci	pci0
pe2	pe	2	0x00000003	L	pci	pci0
pe3	pe	3	0x00000003	L	pci	pci0
ff1	ff	1	0x00000006	SP	pci	quad0
quad1	quad	1	0xff000001	L	sci	sci0
ff2	ff	2	0x00000000	SP	pci	quad1
fabric3	fabric	3	0x00011c00	SMX	fc	ff2
sd713	sd	713	0x2c0118ef	SMX	fc	fabric3
sd714	sd	714	0x2c0118ef	SMX	fc	fabric3
sd715	sd	715	0x2c0118ef	SMX	fc	fabric3
sd716	sd	716	0x2c0118ef	SMX	fc	fabric3
sd713	sd	713	0xac0119ef	SMX	fc	fabric3
sd714	sd	714	0xac0119ef	SMX	fc	fabric3
sd715	sd	715	0xac0119ef	SMX	fc	fabric3
sd716	sd	716	0xac0119ef	SMX	fc	fabric3
ff3	ff	3	0x00000006	SP	pci	quad1

19. Finally, we run an `/etc/diskid` command to see all the devices and the multiple paths which are available to them. Additionally, we can see their vendor specific vital product data, and the size of the LUN, and the detail fabric information on the path.

```
# diskid
cd0: scsibus0 cport 0 target 5 lun 0, No Pbay on scsibus0
  Vendor PLEXTOR, product CD-ROM PX-32CS, revision 1.02
  capacity 45 Mbytes
sd0: scsibus0 cport 4294967295 target 0 lun 0, No Pbay on scsibus0
  Vendor IBM, product DRHS36V, revision 0270, serial # 680196E81K
  capacity 35239 Mbytes
sd713: fabric3 cport 44 fc id 0x118ef, fcp lun 0x1a53
  No Pbay on fabric3
  Vendor IBM, product 2105F20, revision 0.83, serial # 31A20313
  capacity 3814 Mbytes
sd713: fabric3 cport 172 fc id 0x119ef, fcp lun 0x1a53
  No Pbay on fabric3
  Vendor IBM, product 2105F20, revision 0.83, serial # 31A20313
  capacity 3814 Mbytes
sd714: fabric3 cport 44 fc id 0x118ef, fcp lun 0x1b53
  No Pbay on fabric3
```

```

Vendor IBM, product 2105F20, revision 0.83, serial # 31B20313
capacity 3814 Mbytes
sd714: fabric3 cport 172 fc id 0x119ef, fcp lun 0x1b53
No Pbay on fabric3
Vendor IBM, product 2105F20, revision 0.83, serial # 31B20313
capacity 3814 Mbytes
sd715: fabric3 cport 44 fc id 0x118ef, fcp lun 0x1c53
No Pbay on fabric3
Vendor IBM, product 2105F20, revision 0.83, serial # 31C20313
capacity 3814 Mbytes
sd715: fabric3 cport 172 fc id 0x119ef, fcp lun 0x1c53
No Pbay on fabric3
Vendor IBM, product 2105F20, revision 0.83, serial # 31C20313
capacity 3814 Mbytes
sd716: fabric3 cport 44 fc id 0x118ef, fcp lun 0x1d53
No Pbay on fabric3
Vendor IBM, product 2105F20, revision 0.83, serial # 31D20313
capacity 3814 Mbytes
sd716: fabric3 cport 172 fc id 0x119ef, fcp lun 0x1d53
No Pbay on fabric3
Vendor IBM, product 2105F20, revision 0.83, serial # 31D20313
capacity 3814 Mbytes

```

As you can see, there are two entries for each ESS LUN which was assigned to the NUMA-Q system. This is due to the two host ports that were connected to the switch from the ESS. If additional fabrics were added, these devices could be mapped to them also, giving you even higher path redundancy and greater availability. If additional FC HBAs were added to the current fabric from this system, you must add the WWN to the ESS configuration for these devices. Failing to do so would result in the kernel seeing this as an “unbalanced fabric” access to all the devices would be lost.

5.7 Device functionality checks

Now we are ready to checkout our disks to see if we can actually write to and read from them. We recommend the use of the **dd** or **rdcp** for this testing. Examples of both are given here:

```

# dd if=/unix of=/dev/rdisk/sd713 bs=1024 count=25
25+0 records in
25+0 records out
# dd if=/dev/rdisk/sd713 of=/dev/null bs=1024 count=10
10+0 records in
10+0 records out

```

5.8 Specific limitations

The following is a list of some specific limitations which exist within some NUMA-Q environments and not others.

- ▶ With Dynix/ptx V4.4.10 when using the Silkworm 2000 (IBM 2109) type switch, you must ensure that the following switch parameters are turned on.
 - VC-Encoded Address Mode
 - Disable Translative Mode

When set properly, these parameters will show as equal to “1” with a **switchConfigure** command.

- ▶ With Dynix/ptx V4.4.10 the switch domain ID must be set to “0” for the first switch. In a cascaded environment you can have a second switch with a domain ID of “1”. V4.4.10 does not support more than these two domains. V4.4.10 does not support the use of SES (SCSI Enclosure Services).
- ▶ With Dynix/ptx V4.5.x and V4.6.x the switch domain ID is limited by the switch type. These versions of Dynix/ptx do support the SES functions.
 - With Silkworm 1000 type switches domain IDs of 0 - 31 can be used.
 - With the Silkworm 2000 (IBM 2109) type switches domain IDs of 1 - 239 can be used. To avoid conflict with SES default LUN0, do not use the domain ID 0.
- ▶ When using Silkworm 2000 (IBM 2109) type switch in a mixed cascaded environment with a Silkworm 1000 switch, you must ensure that it is configured to run in “compatibility mode”.

For more details on these and other considerations, see the documentation on guidelines for “FC Switches” under “Fibre Channel” at:

<http://webdocs.numaq.ibm.com/>



Compaq AlphaServer Fibre Channel Implementation

In this chapter we describe the steps involved in implementing Fibre Channel host adapter cards in a Compaq AlphaServer running Tru64 UNIX operating system, and its connection to the appropriate SAN switch.

The following topics are discussed:

- ▶ “Introduction” on page 128
- ▶ “Fibre Channel host bus adapters” on page 128
- ▶ “Supported SAN switches” on page 129
- ▶ “Storage” on page 131
- ▶ “Installation procedures” on page 132
- ▶ “Installing the storage system” on page 136

6.1 Introduction

Tru64 UNIX running on Compaq AlphaServer does not support a wide range of Fibre Channel host bus adapters and we recommend that you order the latest or most appropriate one through the Compaq sales channels.

In the topics that follow, we introduce a Compaq 64-bit PCI Fibre Channel host bus adapter for Tru64 UNIX in a step by step installation procedure. In addition to this we cover the supported SAN switches and the considerations which should be taken into account in order to establish the connection with AlphaServer running Tru64 UNIX.

We only discuss a single-path environment, where each fibre device can be accessed via only one path. If you want to implement a multi-path environment, you have to use a storage specific multi-path software such as the IBM Subsystem Device Driver for ESS disk subsystem.

For more general information or technical information on the employment and application of adapters, refer to the Compaq Web site:

<http://www.compaq.com>

6.2 Fibre Channel host bus adapters

In the topics that follow, we describe the Fibre Channel host bus adapters.

6.2.1 Compaq FC HBA

The StorageWorks 64-Bit PCI to Fibre Channel Host Adapter is one of the FC HBAs available from Compaq and supported by AlphaServer running Tru64 UNIX. It belongs to a Compaq KGPSA family, and it can be ordered through P/N 168794-B21 (KGPSA-CA). It replaces the Compaq PCI to Fibre Channel host adapter KGPSA-BC (P/N 380574-001).

As shown in Figure 6-1, this is a single port FC HBA with an embedded multimode (shortwave) optical interface.

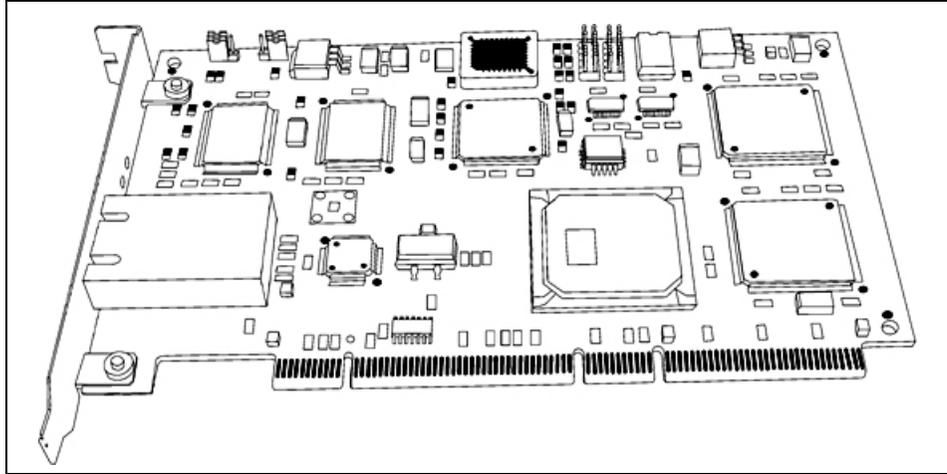


Figure 6-1 Compaq StorageWorks 64-Bit/33MHz PCI-to-FC HBA

The Compaq FC HBA supports:

- ▶ Link bandwidth providing a total throughput of 2 Gb/s
- ▶ Point-To-Point, Arbitrated-Loop and switch fabric FC connections
- ▶ SC type shortwave connection

The number of FC HBAs per server varies and it depends on the AlphaServer model as well as on FC network topology (FC-AL or switch fabric). All the related details can be found in the manual, *64-Bit PCI to Fibre Channel Host Bus Adapter Release Notes for Tru64 UNIX and OpenVMS (AV-RLLUA-TE)*, which can be found at the following Web site:

<http://www.compaq.com/products/storageworks/adapters.html>

The FC HBA drive is already built in Tru64 UNIX operating systems and supported by Console software.

It is recommended to download the latest firmware level in order to ensure that you are at the most current maintenance level.

6.3 Supported SAN switches

The Compaq AlphaServer systems, running the Tru64 UNIX operating system, support only a few SAN switch models, which we cover in this section:

- ▶ IBM 2109-S08
- ▶ IBM 2109-S16
- ▶ Corresponding Brocade 8 and 16 ports switches
- ▶ Corresponding Compaq 8 and 16 ports switches

6.3.1 General information

There are two types of Gigabit Interface Converters (GBIC) and three types of fiber cable commonly used in Fibre Channel environments. FC cabling configurations are described in Table 6-1.

Table 6-1 Common Fibre Channel cabling configurations

GBIC type	Cable type	Distance (1 Gb/s)	Distance (2 Gb/s)
longwave	9 micron, singlemode	10 km	10 km
shortwave	50 micron, multimode	500 m	300 m
shortwave	62.5 micron, multimode	300 m	150 m

Shortwave GBICs and 50 micron cable are usually used to connect devices to the fabric, since most device adapters only support shortwave GBICs, as well as for short inter-switch links. Longwave GBICs and 9 micron cable are used when required for extended distance.

Attenuation

Fibre Channel is a high speed communication technology. Because of this, it is very sensitive to link attenuation. The attenuation is even more of a critical factor when connection speeds rise from 1 Gb/s to 2 Gb/s and even higher.

To limit the effects of attenuation, we recommend making sure that:

- ▶ You have as few connections as possible on a single link.
- ▶ All connectors are clean.
- ▶ All connectors are properly seated.
- ▶ Only one type of fiber is used on a single link (never mix 50 micron and 62.5 micron fiber).

Zoning

We strongly recommend that you use zoning to:

- ▶ Enhance the security of your Fibre Channel fabric.
- ▶ Avoid interactions between different hosts.
- ▶ Secure storage devices that don't support LUN masking, like tape drives.

Since the implementation of zoning and even the terminology used varies between switch manufacturers, we do not discuss it.

6.3.2 IBM 2109-S08/S16

The current supported models of the IBM 2109 family of switches are described in Table 6-2.

Table 6-2 The IBM 2109 family

Model	Ports	Speed	Connectors	Rack space
IBM TotalStorage SAN Switch S08	8	1 Gb/s	SC	1U
IBM TotalStorage SAN Switch S16	16	1 Gb/s	SC	2U

Features and limitations

The IBM 2109 family of switches support both fabric-attached and loop-attached devices on any port, and switch domain ids 1-239 are supported.

Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on customer requirements.

6.4 Storage

The IBM Fibre Channel attached storage systems supported by Compaq AlphaServer running the Tru64 UNIX operating system are listed in Table 6-3.

Table 6-3 IBM Fibre Channel attached disk storage systems

Storage system	HW platforms	Operating systems
IBM TotalStorage 2106 MSS	All AlphaServer models support IBM TotalStorage 2106 MSS	Tru64 UNIX <ul style="list-style-type: none">▶ 4.0F▶ 5.0A

Storage system	HW platforms	Operating systems
IBM TotalStorage 2105 ESS	AlphaServer <ul style="list-style-type: none"> ▶ 800 ▶ 1200 ▶ 2100 ▶ 4000, 4000A ▶ 4100 ▶ 8200, 8400 ▶ DS10, DS20, DS20E ▶ ES40 ▶ GS60, GS60E, GS140 	Tru64 UNIX <ul style="list-style-type: none"> ▶ 4.0F, 4.0G <ul style="list-style-type: none"> – ASE 1.6 ▶ 5.0A <ul style="list-style-type: none"> – TruCluster 5.0A ▶ 5.1 <ul style="list-style-type: none"> – TruCluster 5.1

For additional information related to IBM TotalStorage disk storage devices portfolio, contact your local IBM Storage Specialist.

Other storage system vendors may also support their storage systems on Compaq AlphaServer running Tru64 UNIX.

6.5 Installation procedures

In the following topics, we detail the installation procedures.

6.5.1 Preparing to install the FC HBA

We recommend that at a minimum you check the following:

- ▶ Check if you have a free and supported PCI slot for the FC HBA
 - 32 or 64-bit PCI bus slot
 - Maximum PCI bus clock rate of 33 MHz
- ▶ Tru64 UNIX 4.0F or above operating system level is required
- ▶ Ensure that the console firmware level is v6.0-x or later
- ▶ Download the latest adapter firmware
 - For KGPSA-CA firmware rev. 3.81A4 (2.01A0) or higher
 - For KGPSA-BC firmware rev. 3.01 (1.31) or higher
- ▶ Tru64 patch 399.00 Security (SSR50700U) must be installed

The latest firmware version for the FC HBA can be found at the following Web site:

<http://www.compaq.com/products/storageworks/adapters.html>

6.5.2 Installing the FC HBA

1. Shut down the system.
2. Install the FC HBA in the appropriate PCI slot in the host system.
3. Turn on the system.
4. Bring host system to a halt condition at the console level.
5. Verify the FC HBA installation.
 - Observe the LED indicators on the HBA (see Figure 6-2 and Table 6-4). The green LED indicates power functions and the yellow LED signifies port activity.
 - Normal FC HBA status is indicated by the green LED, which blinks slowly. The yellow one is OFF since there is no connection to either fabrics or storage devices.

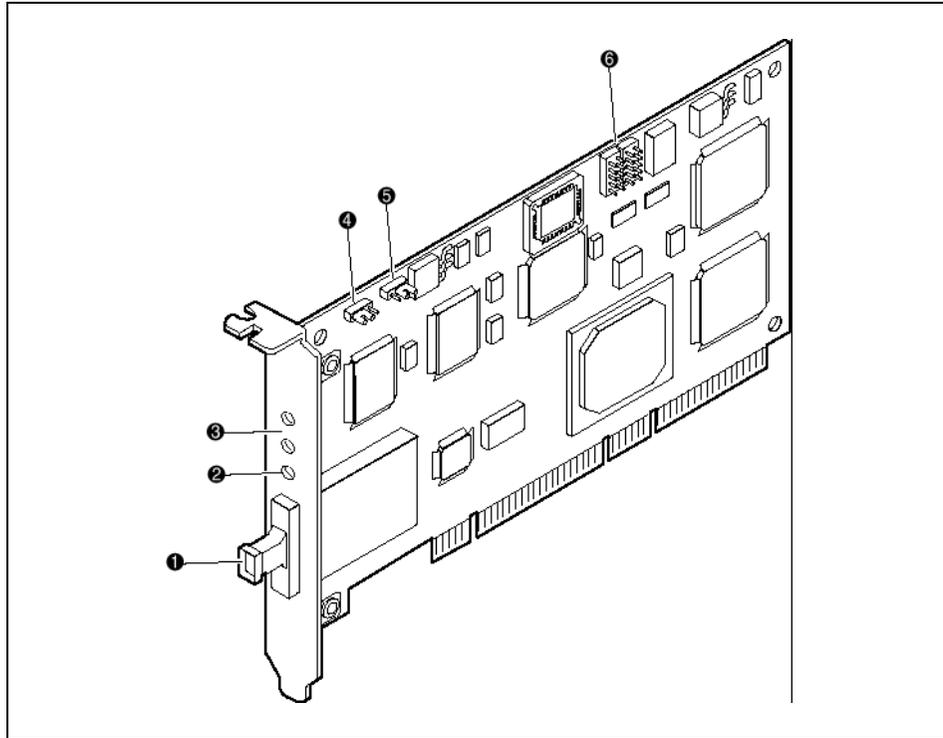


Figure 6-2 64-Bit PCI Fibre Channel Host Adapter

The Figure 6-2 legend is described in Table 6-4.

Table 6-4 64-Bit PCI Fibre Channel HBA description

Figure Legend	Description
1	Fibre Channel connector
2	LED Indicates the Embedded Optical Receive port is receiving light
3	POST LEDs indicators
4 and 5	JUMPERS
6	Not used

6. From the AlphaServer console type:

```
wwidmgr-show adapter
```

This is shown in Example 6-1.

Note: On AlphaServer 8x00, 4x00 or 1200 models, the console must be in diagnostic mode when the **wwidmgr** command is issued. On all other platforms, the **wwidmgr** command can be issued at any time. To get into diagnostic mode, type:

```
set mode diag
```

Example 6-1 Result of set mode diag and wwidmgr -show adapter commands

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

The result of the **wwidmgr -show adapter** command provides basic adapter configuration information such as WWN and SAN topology.

As shown in Example 6-1, **Cur. Topo** (Current Topology), as well as **Next Topo** (Next Topology) value is *FABRIC*. The current topology value can be either *FABRIC* or *LOOP*. It indicates the mode in which the FC HBA is currently operating during this WWIDMGR session.

In the next column, the **Next Topo** value represents FC HBA NVRAM setting and three different states are distinguished: *FABRIC*, *LOOP* or *UNAVAIL*.

In the example just shown in Example 6-1, fabric **Next Topo** value indicates how the FC HBA is configured the next time the system is initialized. If you have a warning message **Nvram Read Failed** and **Next Topo** value is *UNAVAIL*, this means that the FC HBA NVRAM is not formatted. To correct this you should initialize and format FC HBA (as shown in Example 6-2) by use of the **wwidmgr -set adapter** command:

```
wwidmgr -set adapter -item <itemno> -topo <fabric/loop>
```

Example 6-2 Formatting and initializing FC HBA

```
P00>>>wwidmgr -show adapter

item      adapter          WWN              Cur. Topo      Next Topo
kgpsa0.0.0.9.0 -Nvram read failed
[ 0]      kgpsa0.0.0.9.0    2000-0000-c921-df4b  FABRIC        UNAVAIL
kgpsa0.0.0.10.0 -Nvram read failed
[ 1]      kgpsa0.0.0.10.0   2000-0000-c922-69bf  FABRIC        UNAVAIL
[9999] All of the above
P00>>>wwidmgr -set adapter -item 9999 -topo fabric
P00>>>wwidmgr -show adapter

item      adapter          WWN              Cur. Topo      Next Topo
[ 0]      kgpsa0.0.0.9.0    2000-0000-c922-69bf  FABRIC        FABRIC
[ 1]      kgpsa0.0.0.10.0   2000-0000-c922-69bf  FABRIC        FABRIC
[9999] All of the above
P00>>>
```

6.6 Installing the storage system

In this section we describe connection of the Compaq AlphaServer, running Tru64 UNIX, to a storage device after successful FC HBA installation. We connected AlphaServer to a IBM TotalStorage 2105 ESS via IBM 2109-S08 SAN Switch.

The process of storage definition and allocation, as well as switch zones creation are not described here, since they differ between manufacturers. Read carefully the documentation that come with your storage system and your SAN devices.

Switch connectivity check

In order to check if the path from AlphaServer to switch is established, use the **telnet** command to connect to it. As soon as the connection is established, you will be prompted to enter a login name and password. Administrator privileges are necessary in order to make changes in the switch configuration files. The default administrator login name is **admin** and the password is **password**.

At this point you should be familiar with switch interface commands in order to gather switch information or to create a zone. For more information about switch commands and their syntax, type **help** from the switch command line. To check if your FC HBA is properly connected to switch port, type the **switchShow** command, as shown in Example 6-3.

Example 6-3 Switch connectivity check

```
# telnet 9.1.38.58
Trying...
Connected to 9.1.38.158.
Escape character is '^]'.

Fabric OS (tm) Release a2.4.1c

Login: admin
Password:

itsosw2:admin> switchShow
switchName:    itsosw2
switchType:    3.4
switchState:   Online
switchRole:    Subordinate
switchDomain:  2
switchId:      fffc02
switchWwn:     10:00:00:60:69:20:1d:74
switchBeacon:  OFF
port 0: id Online      E-Port 10:00:00:60:69:20:1d:4e "itsosw1" (upstream)
port 1: id No_Light
port 2: -- No_Module
port 3: id No_Light
port 4: id Online      F-Port 50:00:1f:e1:00:0a:e4:02
port 5: id Online      F-Port 10:00:00:00:c9:20:d2:1f
port 6: -- No_Module
port 7: id Online      F-Port 50:06:0b:00:00:07:81:de
itsosw2:admin> quit
Connection closed.
#
```

The online status of the specific port, fabric attachment (F_Port) and FC HBA WWN next to it, are the confirmation that the FC HBA has performed a fabric login to the switch.

Storage connectivity check

At that point it is assumed that volumes on the disk subsystem are created and assigned to AlphaServer.

The following procedures describe how to confirm storage connectivity:

1. Restart the host.
2. Bring the host system to a halt condition at the console level.
3. Type **set mode diag** at the Compaq AlphaServer console (if required by the host) to place the console in diagnostic mode.

4. Type **wwidmgr -show** adapter command to confirm storage attachment as shown in Example 6-4.

Example 6-4 Storage connectivity check

```
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>>>
```

Example 6-4 shows the information about the newly added ESS volumes that you can see at the AlphaServer console. Note that the UDID for each volume appears as -1, signifying that UDID is undefined. The third and fourth quartet of the UDID number is “4942-4d20” and represents a hex string that identifies an ESS volume.

The underscore in Example 6-5 highlights an example of a hex string that identifies the decimal volume number of the ESS volume.

Example 6-5 Decimal volume number for an ESS volume

```
[0] UDID: -1 WWID:01000010:6000-1fe1-4942-4d20-0000-0000-28b1-5660 (ev:none)
```

The last five characters of UDID number are the hex representation of the ESS volume serial number.

In order to verify the attachment of the new added volumes, use the **hwmgr** command as shown in Example 6-6.

Example 6-6 Verifying disk attachment by using hwmgr command

```
# 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
:
:
```

The disk attachment is successful.



Linux on i386 platform Fibre Channel implementation

In this chapter we discuss how to successfully set up a SAN fabric for a Linux server running on i386 platform. We cover the step-by-step procedure from start to finish, and the supported products which work with Linux. We also point out where the latest firmware and device drivers can be obtained.

The following topics are discussed:

- ▶ “Introduction” on page 140
- ▶ “Fibre Channel host adapters” on page 141
- ▶ “SAN fabric” on page 142
- ▶ “Storage” on page 146
- ▶ “Installation procedures” on page 146

7.1 Introduction

Linux is a free, open source operating system that has been gaining popularity in the last few years, especially on the Intel-based platforms. Linux is used extensively as a platform for Internet-based services. In addition to the Intel-based platforms, Linux has been ported to several architectures, including members from all parts of IBM eServer product family.

Unlike other operating systems, there is no single version of Linux available. Several individuals and companies have created *distributions*, usually containing large number of freely available, and in some cases even commercial, applications in addition to the base Linux system. These distributions are usually distributed on CDs, DVDs, or over the Internet. Some of the major distributions available include:

- ▶ Caldera
- ▶ Red Hat
- ▶ SuSE
- ▶ TurboLinux

Because of the number of different distributions available, storage system vendors cannot test and support every possible configuration. Refer to the documentation of your storage system for details on what is supported.

The hardware support for Linux was originally based on the work done by independent developers working at their own pace and time. While the quality of the drivers was generally excellent, the lack of support from the hardware manufacturers meant that it was often hard to find drivers for the very latest hardware. Some hardware manufacturers even decided not to reveal any technical information about their products, making the development of device drivers impossible.

In the last few years, however, many hardware manufacturers have started to realize the market potential Linux has and are developing device drivers to enable their hardware to be used in Linux systems.

7.2 Fibre Channel host adapters

The Fibre Channel host bus adapters supported on Linux by their manufacturers are listed in Table 7-1.

There are also storage system specific limitations on which adapters are supported. Refer to the documentation for your storage system.

Table 7-1 FC HBAs supported in Linux

Adapter	No of ports	Link Speed	Connector
QLogic SANblade 2100 family	1	1 Gb/s	SC
QLogic SANblade 2200 family	1, 2 or 4	1 Gb/s	SC / LC (4port)
QLogic SANblade 2300 family	1 or 2	1 or 2 Gb/s	LC
Emulex LP7000E	1	1 Gb/s	SC
Emulex LP8000	1	1 Gb/s	SC
Emulex LP8000DC	2	1 Gb/s	SC
Emulex LP850	1	1 Gb/s	SC
Emulex LP9000	1	1 Gb/s	SC
Emulex LP9002L	1	1 or 2 Gb/s	LC
Emulex LP9002DC	2	1 or 2 Gb/s	LC
Emulex LP9402DC	2	1 or 2 Gb/s	LC
Emulex LP952L	2	1 or 2 Gb/s	LC
JNI FCE-6410	1	1 Gb/s	SC
JNI FCE-3210	1	1 Gb/s	SC
JNI FCE2-6410	2	1 Gb/s	SC
JNI FCE2-6412	2	1 Gb/s	SC

7.2.1 QLogic

The current QLogic driver for Linux supports Linux kernel versions 2.2.16 and 2.4.x. The complete driver is distributed as source code.

You can find more information on QLogic adapters, as well as drivers for them at:

<http://www.qlogic.com/>

7.2.2 Emulex

The current Emulex driver for Linux supports Linux kernel versions 2.2.19, 2.4.2, 2.4.3, and 2.4.4. It also supports both FCP and IP protocols on a single Fibre Channel adapter simultaneously. We do not discuss the IP functionality.

You can find more information on Emulex adapters, as well as drivers for them, at:

<http://www.emulex.com/>

7.2.3 JNI

The current JNI drivers for Linux supports only Linux kernel versions 2.2.12 and 2.2.14. Since the drivers are only distributed in binary format, it is not possible to get the current drivers to work on current 2.4.x kernels. For that reason, we do not test installation of JNI adapters.

You can find more information on JNI adapters, as well as drivers for them, at:

<http://www.jni.com/>

7.3 SAN fabric

Linux supports three Fibre Channel network topologies, depending on the Fibre Channel adapter:

- ▶ Fibre Channel Arbitrated Loop (FC-AL)
- ▶ Fibre Channel point-to-point
- ▶ Switched fabric

FC-AL has severe performance and scalability limitations. Fibre Channel point-to-point is usually not feasible either, because of the high cost and low number of Fibre Channel ports in any storage system. Because of these limitations, we only cover switched fabric topology. However, we also discuss attaching storage devices supporting only FC-AL attachment to the fabrics.

We cover the products from the following three vendors:

- ▶ IBM 2109 family
- ▶ INRANGE
- ▶ McDATA

7.3.1 General information

There are two types of Gigabit Interface Converters (GBIC) and three types of fiber cable commonly used in Fibre Channel environments, as described in Table 7-2.

Table 7-2 Common Fibre Channel cabling configurations

GBIC type	Cable type	Distance (1 Gb/s)	Distance (2 Gb/s)
longwave	9 micron, singlemode	10 km	10 km
shortwave	50 micron, multimode	500 m	300 m
shortwave	62.5 micron, multimode	300 m	150 m

Shortwave GBICs and 50 micron cable are usually used to connect devices to the fabric, since most device adapters only support shortwave GBICs, as well as for short inter-switch links. Longwave GBICs and 9 micron cable are used when required for extended distance.

Attenuation

Fibre Channel is a high speed communication technology. Because of this, it is very sensitive to link attenuation. The attenuation is even more of a critical factor when connection speeds rise from 1 Gb/s to 2 Gb/s and even higher.

To limit the effects of attenuation, we recommend making sure that:

- ▶ You have as few connections as possible on a single link.
- ▶ All connectors are clean.
- ▶ All connectors are properly seated.
- ▶ Only one type of fiber is used on a single link (never mix 50 micron and 62.5 micron fiber).

Zoning

We strongly recommend that you use zoning to:

- ▶ Enhance the security of your Fibre Channel fabric.
- ▶ Avoid interactions between different hosts.
- ▶ Secure storage devices that do not support LUN masking, like tape drives.

Since the implementation of zoning and even the terminology used varies between switch manufacturers, we do not discuss it in any detail.

7.3.2 IBM 2109 family

The current members of the IBM 2109 family of switches are described in Table 7-3.

Table 7-3 The IBM 2109 family

Model	Ports	Speed	Connectors	Rack space
IBM TotalStorage SAN Switch S08	8	1 Gb/s	SC	1U
IBM TotalStorage SAN Switch S16	16	1 Gb/s	SC	2U
IBM TotalStorage SAN Switch F16	16	1 or 2 Gb/s	LC	1U

Features and limitations

The IBM 2109 family of switches support both fabric-attached and loop-attached devices on any port, and switch domain ids 1-239 are supported.

Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on your requirements.

For more information, see *IBM SAN Survival Guide Featuring the IBM 2109*, SG24-6127.

7.3.3 INRANGE

The INRANGE switches currently supported by IBM (the IBM 2042 family) are described in Table 7-4.

Table 7-4 The INRANGE switches

Model	Ports	Speed	Connectors	Rack space
2042-001 INRANGE FC/9000 Fibre Channel Director	64	1 Gb/s	SC	20U
2042-128 INRANGE FC/9000-128 Fibre Channel Director	128	1 Gb/s	SC	40U

In addition, IBM also sells the INRANGE FC/9000 40U cabinet (2042-C40). The cabinet is able to hold one 128-port switch or two 64-port switches.

Features and limitations

The INRANGE FC/9000 family of switches support both fabric-attached and loop-attached devices on any port.

Any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC, depending on your requirements.

For more information, see *IBM SAN Survival Guide Featuring the INRANGE Portfolio*, SG24-6150.

7.3.4 McDATA

The McDATA switches currently supported by IBM (the IBM 2031 and 2032 families) are described in Table 7-5.

Table 7-5 The McDATA switches

Model	Ports	Speed	Connectors	Rack space
2031-016 McDATA ES-3016 Fabric Switch	16	1 Gb/s	LC	1U
2031-032 McDATA ES-3032 Fabric Switch	32	1 Gb/s	LC	1.5U
2031-L00 McDATA ES-1000 Loop Switch	8+1	1 Gb/s	SC	1U
2032-001 McDATA ED-5000 Enterprise Fibre Channel Director	32	1 Gb/s	SC	18U
2032-064 McDATA ED-6064 Enterprise Fibre Channel Director	64	1 Gb/s	LC	9U

In addition, IBM also sells the McDATA FC-512 Cabinet (2032-C36).

Features and limitations

The only McDATA switch that supports loop attached devices (tape drives, for example) is the McDATA ES-1000 Loop Switch.

The ports in the 2032-001 and 2032-064 can be added in cards. The cards are not compatible between the models. Each card contains four ports. There are three different cards available:

- ▶ Four shortwave (multimode) ports
- ▶ Four longwave (singlemode) ports
- ▶ Three shortwave ports and one longwave port

On the 2031-016 and 2031-032 switches, any port can be independently configured with a shortwave (multimode) or longwave (singlemode) GBIC.

The 2031-L00 switch has two different types of ports:

- ▶ Eight switching Hub ports that can only be configured with shortwave (multimode) GBIC (or copper interface)
- ▶ One bridge port, for connecting the loop switch to the fabric, that can be configured with a shortwave (multimode) or longwave (singlemode) GBIC.

The McDATA switches support fabrics of up to 31 switches only.

For more information, see *IBM SAN Survival Guide Featuring the McDATA Portfolio*, SG24-6149.

7.4 Storage

The IBM Fibre Channel attached disk storage systems supported on Linux systems are described in Table 7-6. There is no support available for IBM Fibre Channel attached tape drives on Linux at the moment.

Table 7-6 IBM Fibre Channel attached disk storage systems

Storage system	Supported distributions	Notes
IBM TotalStorage Enterprise Storage Server	Red Hat 7.1 SuSE 7.2	
IBM TotalStorage FAStT700 Storage Server	Red Hat 7.1	Only supports IBM 2109 switch fabrics
IBM TotalStorage FAStT500 Storage Server	Red Hat 7.1	
IBM TotalStorage FAStT200 Storage Server	Red Hat 7.1	Only supports IBM 2109 switch fabrics

Other storage system vendors may also support their storage systems on Linux systems.

7.5 Installation procedures

Installing new devices in a Linux system can be quite different from installing the devices in any other operating system. The device drivers are often delivered as source files, written usually in the C programming language, and have to be compiled before they can be used.

There are two ways to install a device driver in Linux. Device drivers can either be compiled directly into the kernel (“monolithic” kernel), or they can be compiled as loadable kernel modules. Most device drivers are compatible with both types of installation. We chose to install the drivers as modules, since that gives us much more flexibility.

We choose Red Hat Linux 7.1 as the starting point of our installations, because that distribution is currently supported by the widest range of FC HBAs and storage devices. The installation procedure is, however, similar in most other Linux distributions.

7.5.1 Installing a QLogic FC HBA

We use the following in our installation:

- ▶ IBM Netfinity 5500
- ▶ Red Hat Linux 7.1, with kernel sources and kernel includes installed
- ▶ Kernel 2.4.2
- ▶ QLogic QLA2200F FC HBA
- ▶ QLogic driver 4.27beta

Red Hat Linux already contains modules for QLogic adapters, but because it is an old version, we rebuild a new one from the sources.

Our installation procedure is as follows:

1. Start up your system and log in as root.
2. Check that the kernel source files and headers are installed, as follows:

```
[root@clyde /root]# rpm -q -a |grep kernel
kernel-headers-2.4.2-2
kernel-2.4.2-2
kernel-source-2.4.2-2
[root@clyde /root]#
```

You need to have both kernel-headers and kernel-source installed. If either of these packages is missing, you have to install it before proceeding.

3. Add the following line to the file /etc/lilo.conf:

```
append="max_scsi_luns=128"
```

If you already have an append line in /etc/lilo.conf, add the max_scsi_luns parameter to that line.

4. Reinstall LiLO by the following command:

```
[root@clyde /root]# /sbin/lilo
Added linux *
[root@clyde /root]#
```

5. Create a directory for the driver package and unpack it there, as follows:

```
[root@clyde /root]# mkdir /usr/src/qla2x00
[root@clyde /root]# cd /usr/src/qla2x00
[root@clyde qla2x00]# tar xvvfz /tmp/qla2x00src-v4\[1\].27Beta.tgz
-rw-r--r-- root/root      1924 2001-04-10 10:13:42 AddKernel.txt
-rw-r--r-- root/root      3205 2001-04-17 12:43:07 Makefile
-rw-r--r-- root/root     10887 2001-04-23 19:31:55 README.qla2x00
-rw-r--r-- root/root     25469 2001-03-09 15:58:43 exioct.h
-rw-r--r-- root/root      4149 2001-03-09 15:58:43 exioctln.h
-rw-r--r-- root/root    306034 2001-03-09 15:58:43 ql2100_fw.h
-rw-r--r-- root/root    327593 2001-03-09 15:58:43 ql2200_fw.h
-rw-r--r-- root/root    338250 2001-03-09 15:58:43 ql2300_fw.h
-rw-r--r-- root/root    524896 2001-04-16 18:40:30 qla2x00.c
-rw-r--r-- root/root      87171 2001-04-16 16:56:47 qla2x00.h
-rw-r--r-- root/root     49375 2001-03-09 15:58:43 qla2x00ioctl.c
-rw-r--r-- root/root     16993 2001-03-09 15:58:43 qlavendor.c
-rw-r--r-- root/root      3320 2001-04-23 19:31:56 release.txt
[root@clyde qla2x00]#
```

6. Build the **qla2x00.o** loadable kernel module:

```
[root@clyde qla2x00]# make OSVER=linux-2.4
cc -D__KERNEL__ -DMODULE -Wall -O2 -DISP2200 -DMODVERSIONS -include
/usr/src/linux-2.4/include/linux/modversions.h -I/usr/src/linux-2.4/include
-I/usr/src/linux-2.4/include/./drivers/scsi -Wall -Wstrict-prototypes
-fomit-frame-pointer -fno-strength-reduce -pipe -m486 -malign-loops=2 -malign-jumps=2
-malign-functions=2 -DCONFIG_X86_LOCAL_APIC -c -o qla2x00.o qla2x00.c
[root@clyde qla2x00]#
```

Note: This procedure assumes that you are using a single processor system.

7. Copy the new module into your modules directory:

```
[root@clyde qla2x00]# cp qla2x00.o /lib/modules/2.4.2-2/kernel/drivers/scsi/
cp: overwrite `/lib/modules/2.4.2-2/kernel/drivers/scsi/qla2x00.o'? y
[root@clyde qla2x00]#
```

8. Shut down your system.
9. Install the QLogic adapter.
10. Connect the adapter to the Fibre Channel network.
11. Start up the system and log in as root.

Note: During the startup, Linux may notice that you have installed the adapter and want to configure it. Let Linux configure the adapter by selecting **Configure**.

12. Rebuild the module dependencies file:

```
[root@clyde /root]# depmod -a
[root@clyde /root]#
```

13. Load the module as follows:

```
[root@clyde /root]# modprobe qla2x00
[root@clyde /root]#
```

The command may take some time to complete.

14. Check the kernel messages with `dmesg` command, as in Example 7-1.

Example 7-1 Checking the kernel messages from loading of QLogic driver

```
[root@clyde /root]# dmesg
...
qla2x00: detect() found an HBA
qla2x00: VID=1077 DID=2200 SSVID=1077 SSDID=2
(scsi): Found a QLA2200 @ bus 1, device 0x4, irq 11, iobase 0x5200
scsi(2): Configure NVRAM parameters...
scsi(2): Verifying loaded RISC code...
scsi(2): Verifying chip...
scsi(2): LIP reset occurred
scsi(2): Waiting for LIP to complete...
scsi(2): LOOP UP detected
spurious 8259A interrupt: IRQ7.
qla2100: Performing ISP error recovery - ha= f70e0078
scsi(2): LIP reset occurred
scsi(2): Waiting for LIP to complete...
scsi(2): Waiting for LIP to complete...
qla2100_configure_hba: Get host loop ID failed
scsi-qla0-adapter-node=200000e08b0190ef;
scsi-qla0-adapter-port=210000e08b0190ef;
scsi2 : QLogic QLA2200 PCI to Fibre Channel Host Adapter: bus 1 device 4 irq 11
        Firmware version: 2.01.27, Driver version 4.27b
scsi(2): LOOP UP detected
scsi(2): Waiting for LIP to complete...
scsi2: Topology - (F_Port), Host Loop address 0xffff
scsi(2): Waiting for LIP to complete...
scsi2: Topology - (F_Port), Host Loop address 0xffff
[root@clyde /root]#
```

15. You can also display information about the FC HBA from `/proc` filesystem, as in Example 7-2.

Example 7-2 Checking the QLogic adapter via /proc interface

```
[root@clyde /root]# cat /proc/scsi/qla2x00/2
QLogic PCI to Fibre Channel Host Adapter for ISP2100/ISP2200/ISP2300:
    Firmware version: 2.01.27, Driver version 4.27b
Request Queue = 0x370f8000, Response Queue = 0x37100000
Request Queue count= 256, Response Queue count= 256
Number of pending commands = 0x0
Number of queued commands = 0x0
Number of free request entries = 253
Number of mailbox timeouts = 1
Number of ISP aborts = 1
Number of loop resyncs = 1
Number of retries for empty slots = 0
Number of reqs in retry_q = 0
Number of reqs in done_q = 0

SCSI Device Information:
scsi-qla0-adapter-node=200000e08b0190ef;
scsi-qla0-adapter-port=210000e08b0190ef;
scsi-qla0-target-0=50001fe1000ae400;
scsi-qla0-target-1=50001fe1000ae400;

SCSI LUN Information:
(Id:Lun)
[root@clyde /root]#
```

You can see the WWN of your adapter as the parameter `scsi-qla0-adapter-port`. You will need that parameter to configure your storage system. Note also that you do not see any SCSI LUNs yet.

16. Optionally, check to see that the port is correctly configured on the switch.

The procedure you need to use is different for each switch vendor. We show the procedure for the IBM 2109 family of switches in Example 7-3. For INRANGE switches, you should use the IN-VSN Enterprise Manager application to find the same information. For McDATA switches, you can find the information in the EFC Manager application.

Example 7-3 Checking QLogic adapter switch connectivity

```
[root@clyde /root]# telnet 9.1.38.159
Trying 9.1.38.159...
Connected to 9.1.38.159.
Escape character is '^'.
```

```
Fabric OS (tm) Release a2.4.1c
```

```
login: admin
```

```
Password:
```

```

itsosw3:admin> switchShow
switchName:    itsosw3
switchType:    2.4
switchState:   Online
switchRole:    Principal
switchDomain:  3
switchId:      fffc03
switchWwn:     10:00:00:60:69:10:64:cf
switchBeacon:  OFF
port 0: id Online      F-Port  21:00:00:e0:8b:02:a7:2d
port 1: -- No_Module
port 2: id Online      E-Port  10:00:00:60:69:20:1d:4e "itsosw1" (downstream)
port 3: -- No_Module
port 4: id No_Light
port 5: id Online      F-Port  21:00:00:e0:8b:01:90:ef
port 6: id Online      E-Port  10:00:00:60:69:20:1d:4e "itsosw1"
port 7: id No_Light
port 8: sw Online      F-Port  10:00:00:00:c9:22:bf:75
port 9: -- No_Module
port 10: sw No_Light
port 11: sw No_Light
port 12: sw Online     L-Port  1 public
port 13: id No_Light
port 14: id Online     F-Port  10:00:00:00:c9:25:8e:54
port 15: id No_Light
itsosw3:admin> exit
Connection closed by foreign host.
[root@clyde /root]#

```

From Example 7-3, we see that the port we are using (port 5) is online, it is an F_Port (for fabric attachment), and the WWN listed is the same that we see on our adapter in Example 7-2.

17. If you want your FC HBA to be activated automatically in system startup, add the following lines to `/etc/rc.d/rc.sysinit` before the Add raid devices comment:

```

# Add Fibre Channel devices
modprobe scsi_hostadapter1

```

18. Your FC HBA is now correctly installed.

7.5.2 Configuring storage for the QLogic FC HBA

After the adapter installation, we have to create any necessary switch zones and allocate the storage for our host from the storage system we are using. As both of those steps are different for different switch manufacturers and different storage subsystems, they are not described here.

We have access to a single volume in a storage subsystem. We want to configure that volume to our host system. The procedure is as follows:

1. Remove and reinstall the QLogic driver as follows:

```
[root@clyde /root]# rmmod qla2x00
[root@clyde /root]# modprobe qla2x00
[root@clyde /root]#
```

Alternatively, you can reboot the machine.

2. Check the current information on the FC HBA from /proc filesystem, as in Example 7-4.

Example 7-4 Checking the QLogic adapter again via /proc interface

```
[root@clyde /proc]# cat /proc/scsi/qla2x00/2
QLogic PCI to Fibre Channel Host Adapter for ISP2100/ISP2200/ISP2300:
    Firmware version: 2.01.27, Driver version 4.27b
Request Queue = 0x370f8000, Response Queue = 0x37040000
Request Queue count= 256, Response Queue count= 256
Number of pending commands = 0x0
Number of queued commands = 0x0
Number of free request entries = 245
Number of mailbox timeouts = 0
Number of ISP aborts = 0
Number of loop resyncs = 0
Number of retries for empty slots = 0
Number of reqs in retry_q = 0
Number of reqs in done_q = 0

SCSI Device Information:
scsi-qla0-adapter-node=200000e08b0190ef;
scsi-qla0-adapter-port=210000e08b0190ef;
scsi-qla0-target-0=50001fe1000ae400;
scsi-qla0-target-1=50001fe1000ae400;

SCSI LUN Information:
(Id:Lun)
( 0: 0): Total reqs 5, Pending reqs 0, Queued reqs 0,
[root@clyde /proc]#
```

Note that now you also see a LUN behind the FC HBA.

3. Check the kernel messages with the **dmesg** command, as in Example 7-5.

Example 7-5 Checking the kernel messages

```
[root@clyde /root]# dmesg
...
qla2x00: detect() found an HBA
qla2x00: VID=1077 DID=2200 SSVID=1077 SSDID=2
(scsi): Found a QLA2200 @ bus 1, device 0x4, irq 11, iobase 0x5200
scsi(2): Configure NVRAM parameters...
```

```

scsi(2): Verifying loaded RISC code...
scsi(2): Verifying chip...
scsi(2): Waiting for LIP to complete...
scsi(2): LOOP UP detected
scsi(2): Waiting for LIP to complete...
scsi2: Topology - (F_Port), Host Loop address 0xffff
scsi(2): Waiting for LIP to complete...
scsi2: Topology - (F_Port), Host Loop address 0xffff
scsi-qla0-adapter-node=200000e08b0190ef;
scsi-qla0-adapter-port=210000e08b0190ef;
scsi-qla0-target-0=50001fe1000ae400;
scsi-qla0-target-1=50001fe1000ae400;
scsi2 : QLogic QLA2200 PCI to Fibre Channel Host Adapter: bus 1 device 4 irq 11
        Firmware version: 2.01.27, Driver version 4.27b
        Vendor: DEC          Model: HSG80          Rev: V85S
        Type:   Direct-Access      ANSI SCSI revision: 02
        Vendor: DEC          Model: HSG80          Rev: V85S
        Type:   Processor         ANSI SCSI revision: 02
scsi(2:0:0:0): Enabled tagged queuing, queue depth 16.
scsi(2:0:1:0): Enabled tagged queuing, queue depth 16.
Attached scsi disk sdb at scsi2, channel 0, id 0, lun 0
SCSI device sdb: 17418235 512-byte hdwr sectors (8918 MB)
sdb:
[root@clyde /root]#

```

Note that now we also see the ID of our disk (sdb). If you already had partitions on the disk, they would also be shown.

4. Start using the new LUN like any normal SCSI disk.

We tested the Fibre Channel drive according to the procedure in Example 7-6.

Example 7-6 Testing a Fibre Channel disk with QLogic adapter

```

[root@clyde /root]# fdisk /dev/sdb
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel. Changes will remain in memory only,
until you decide to write them. After that, of course, the previous
content won't be recoverable.

```

The number of cylinders for this disk is set to 8504.
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): **p**

Disk /dev/sdb: 64 heads, 32 sectors, 8504 cylinders

Units = cylinders of 2048 * 512 bytes

```
Device Boot Start End Blocks Id System
```

Command (m for help): **n**

Command action

e extended

p primary partition (1-4)

p

Partition number (1-4): **1**

First cylinder (1-8504, default 1): **1**

Last cylinder or +size or +sizeM or +sizeK (1-8504, default 8504): **8504**

Command (m for help): **w**

The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: If you have created or modified any DOS 6.x partitions, please see the fdisk manual page for additional information.

Syncing disks.

```
[root@clyde /root]# mkfs /dev/sdb1
```

mke2fs 1.19, 13-Jul-2000 for EXT2 FS 0.5b, 95/08/09

Filesystem label=

OS type: Linux

Fragment size=4096 (log=2)

1089152 inodes, 2177020 blocks

108851 blocks (5.00%) reserved for the super user

First data block=0

67 block groups

32768 blocks per group, 32768 fragments per group

16256 inodes per group

Superblock backups stored on blocks:

32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Writing inode tables: done

Writing superblocks and filesystem accounting information: done

```
[root@clyde /root]# mkdir /test
```

```
[root@clyde /root]# mount /dev/sdb1 /test
```

```
[root@clyde /root]# df -k
```

Filesystem	1k-blocks	Used	Available	Use%	Mounted on
/dev/sda7	402885	84463	297621	23%	/
/dev/sda2	201433	3822	187210	3%	/boot
/dev/sda5	3024016	538392	2332008	19%	/usr
/dev/sdb1	8571328	20	8135904	1%	/test

```
[root@clyde /root]# cp /etc/hosts /test
```

```
[root@clyde /root]# ls -l /etc/hosts /test/hosts
```

```
-rw-r--r-- 1 root root 181 Dec 3 04:00 /etc/hosts
```

```
-rw-r--r-- 1 root root 181 Dec 3 15:15 /test/hosts
```

```
[root@clyde /root]# diff /etc/hosts /test/hosts
```

```
[root@clyde /root]#
```

The connection is successful.

7.5.3 Installing an Emulex FC HBA

We used the following in our installation:

- ▶ IBM Netfinity 5500
- ▶ Red Hat Linux 7.1, with kernel sources and kernel includes installed
- ▶ Kernel 2.4.2
- ▶ Emulex LP8000 FC HBA
- ▶ Emulex driver SCSI/IP v4.11e

Our installation procedure is as follows:

1. Start up your system and log in as root.
2. Check that the kernel source files and headers are installed, as follows:

```
[root@clyde /root]# rpm -q -a |grep kernel
kernel-headers-2.4.2-2
kernel-2.4.2-2
kernel-source-2.4.2-2
[root@clyde /root]#
```

You need to have both kernel-headers and kernel-source installed. If either of these packages is missing, you have to install it before proceeding.

3. Add the following line to the file /etc/lilo.conf:

```
append="max_scsi_luns=128"
```

If you already have an append line in /etc/lilo.conf, add the max_scsi_luns parameter to that line.

4. Reinstall LiLO by the following command:

```
[root@clyde /root]# /sbin/lilo
Added linux *
[root@clyde /root]#
```

5. Create a directory for the driver package and unpack it there, as follows:

```
[root@clyde /root]# mkdir /usr/src/emulex
[root@clyde /root]# cd /usr/src/emulex
[root@clyde emulex]# tar xvvf /tmp/lpfc-i386.tar
drwxr-xr-x root/root      0 2001-10-01 17:09:14 ./SourceBuild/
drwxr-xr-x root/root      0 2001-09-27 12:01:10 ./SourceBuild/include/
-rw-r--r-- root/root    27645 2001-09-27 12:01:10 ./SourceBuild/include/fc_os.h
-rw-r--r-- root/root     5662 2001-09-27 12:01:10 ./SourceBuild/include/fc_ertn.h
-rw-r--r-- root/root     8548 2001-09-27 12:01:10 ./SourceBuild/include/dfc.h
-rw-r--r-- root/root    75966 2001-09-27 12:01:10 ./SourceBuild/include/fc.h
-rw-r--r-- root/root    95934 2001-09-27 12:01:10 ./SourceBuild/include/fc_hw.h
-rw-r--r-- root/root    12439 2001-09-27 12:01:10 ./SourceBuild/include/fcdds.h
-rw-r--r-- root/root    16251 2001-09-27 12:01:10 ./SourceBuild/include/fcdiag.h
-rw-r--r-- root/root    18532 2001-09-27 12:01:10 ./SourceBuild/include/fcgparm.h
-rw-r--r-- root/root    15720 2001-09-27 12:01:10 ./SourceBuild/include/fc_crtn.h
-rw-r--r-- root/root    177693 2001-09-27 12:01:10 ./SourceBuild/lpfdriver
-rw-r--r-- root/root     16993 2001-09-27 12:01:10 ./SourceBuild/lpfc.conf.c
```

```

-rw-r--r-- root/root      186117 2001-10-01 12:43:37 ./SourceBuild/fcLINUXfcp.c
-r--r--r-- root/root      3242 2001-09-27 12:01:10 ./SourceBuild/Makefile
-r--r--r-- root/root      3242 2001-09-27 12:01:10 ./SourceBuild/Makefile.module.i386
-r--r--r-- root/root      3280 2001-09-27 12:01:10 ./SourceBuild/Makefile.kernel
-r--r--r-- root/root       384 2001-09-27 12:01:10 ./SourceBuild/diag_load
-r--r--r-- root/root       163 2001-09-27 12:01:10 ./SourceBuild/diag_unload
-rw-r--r-- root/root      3627 2001-10-01 17:09:14 ./SourceBuild/README
-rw-r--r-- root/root       643 2001-09-27 12:01:10 ./SourceBuild/Install.sh
-rw-r--r-- root/root       378 2001-09-27 12:01:10 ./SourceBuild/Remove.sh
-rwxr-xr-x root/root      83818 2001-09-27 12:01:10 ./SourceBuild/dfc
-rw-r--r-- root/root     13978 2001-09-27 12:01:10 ./SourceBuild/libdfc.a
-rwxr-xr-x root/root     102514 2001-09-27 12:01:10 ./SourceBuild/lputil
-rwxr-xr-x root/root     14340 2001-09-27 12:01:10 ./SourceBuild/libHBAAPI.so
-rwxr-xr-x root/root     77946 2001-09-27 12:01:10 ./SourceBuild/libemulexhaapi.so
[root@clyde emulex]# cd SourceBuild/
[root@clyde SourceBuild]#

```

6. Change the parameters for the Makefile according to your system:

- If your kernel source resides in any other directory than /usr/src/linux, change the BASEINCLUDE variable to point to the correct directory.
- If you are running a single-processor kernel, remove the line `-D__SMP__=1 \` from the CFLAGS variable.

7. Check the parameter `lpfc_topology` in the file `lpfc.conf.c` and change it, if necessary.

Note: By default, the value of `lpfc_topology` is set to 4. That means, that the adapter will only try loop-mode login and not point-to-point mode. We change the value to 6, meaning that the adapter will try first point-to-point mode, and then loop mode, if point-to-point mode fails.

8. Build the `lpfcdd` and `lpfndd` loadable kernel modules:

```

[root@clyde SourceBuild]# make build
Build Environment root: /usr/src/linux-2.4
cc -D__GENKSYMS__ -D__KERNEL__=1 -DMODULE -DMODVERSIONS -include
/usr/src/linux-2.4/include/linux/modversions.h -I./include
-I/usr/src/linux-2.4/drivers/scsi -I/usr/src/linux-2.4/include/scsi
-I/usr/src/linux-2.4/include -DLP6000 -D_LINUX -I./include
-I/usr/src/linux-2.4/drivers/scsi -I/usr/src/linux-2.4/include/scsi
-I/usr/src/linux-2.4/include -E fcLINUXfcp.c > lpfc.ver1
cat lpfc.ver1 | /sbin/genksyms -k 2.2.5 > lpfc.ver
cc -Wall -O2 -fomit-frame-pointer -D__KERNEL__=1 -DMODULE -DMODVERSIONS -include
/usr/src/linux-2.4/include/linux/modversions.h -I./include
-I/usr/src/linux-2.4/drivers/scsi -I/usr/src/linux-2.4/include/scsi
-I/usr/src/linux-2.4/include -DLP6000 -D_LINUX -I./include
-I/usr/src/linux-2.4/drivers/scsi -I/usr/src/linux-2.4/include/scsi
-I/usr/src/linux-2.4/include -c fcLINUXfcp.c
rm -f lpfc.ver1

```

```

cc -Wall -O2 -fomit-frame-pointer -D__KERNEL__=1 -DMODULE -DMODVERSIONS -include
/usr/src/linux-2.4/include/linux/modversions.h -I./include
-I/usr/src/linux-2.4/drivers/scsi -I/usr/src/linux-2.4/include/scsi
-I/usr/src/linux-2.4/include -DLP6000 -D_LINUX -I./include
-I/usr/src/linux-2.4/drivers/scsi -I/usr/src/linux-2.4/include/scsi
-I/usr/src/linux-2.4/include -c fcLINUXlan.c
rm -f lpfc.ver
cc -Wall -O2 -fomit-frame-pointer -D__KERNEL__=1 -DMODULE -DMODVERSIONS -include
/usr/src/linux-2.4/include/linux/modversions.h -I./include
-I/usr/src/linux-2.4/drivers/scsi -I/usr/src/linux-2.4/include/scsi
-I/usr/src/linux-2.4/include -DLP6000 -D_LINUX -I./include
-I/usr/src/linux-2.4/drivers/scsi -I/usr/src/linux-2.4/include/scsi
-I/usr/src/linux-2.4/include -c lpfc.conf.c
cp lpfcdriver lpfcdriver.o
ld -r -o lpfcdd.2.4.2-2.o lpfcdriver.o fcLINUXfcp.o lpfc.conf.o
ld -r -o lpfndd.2.4.2-2.o fcLINUXlan.o
[root@clyde SourceBuild]#

```

9. Install the new module into your modules directory:

```

[root@clyde SourceBuild]# make install
Installing under /lib/modules/2.4.2-2/kernel/drivers
[root@clyde SourceBuild]#

```

10. Copy the Emulex tools **dfc** and **lputil** to **/usr/sbin**:

```

[root@clyde SourceBuild]# cp dfc lputil /usr/sbin
[root@clyde SourceBuild]#

```

11. Add the following line to **/etc/modules.conf** to configure the new module:

```

alias scsi_hostadapter1 lpfcdd

```

12. Rebuild the module dependencies file:

```

[root@clyde /root]# depmod -a
[root@clyde /root]#

```

13. Shut down your system.

14. Install the Emulex adapter.

15. Connect the adapter to the Fibre Channel network.

16. Start up the system and log in as root.

17. Load the module as follows:

```

[root@clyde /root]# modprobe lpfcdd
[root@clyde /root]#

```

18. Check the kernel messages with the **dmesg** command:

```

[root@clyde /root]# dmesg
...
Emulex LightPulse FC SCSI/IP 4.11e
!lpfc-1962934273:045:Vital Product Data Data: 82 23 0 36
!lpfc-1962934273:031:Link Up Event received Data: 1 1 1 0

```

```
scsi2 : Emulex LPFC (DRAGONFLY) SCSI on PCI bus 01 device 20 irq 11
[root@clyde /root]#
```

19. Check the adapter with the **lputil** command, as in Example 7-7.

Example 7-7 Checking the Emulex adapter with the lputil command

```
[root@clyde /root]# lputil
LightPulse Common Utility for Linux. Version 1.2a22 (5/17/2001).
Copyright (c) 2000, Emulex Network Systems, Inc.
```

```
Emulex Fibre Channel Host Adapters Detected: 1
Host Adapter 0 (lpfc-1962934273) is an LP8K_V2 (Ready Mode)
```

MAIN MENU

1. List Adapters
2. Adapter Information
3. Firmware Maintenance
4. Reset Adapter

0. Exit

Enter choice => **2**

ADAPTER INFORMATION MENU

1. PCI Configuration Parameters
2. Adapter Revision Levels
3. Wakeup Parameters
4. IEEE Address
5. Loop Map
6. Status & Counters
7. Link Status
8. Configuration Parameters

0. Return to Main Menu

Enter choice => **4**

IEEE Address for Adapter 0:

```
[10000000] [C923C305]
```

ADAPTER INFORMATION MENU

1. PCI Configuration Parameters
2. Adapter Revision Levels
3. Wakeup Parameters
4. IEEE Address

5. Loop Map
 6. Status & Counters
 7. Link Status
 8. Configuration Parameters
-
0. Return to Main Menu

Enter choice => **0**

MAIN MENU

1. List Adapters
 2. Adapter Information
 3. Firmware Maintenance
 4. Reset Adapter
-
0. Exit

Enter choice => **0**

[root@clyde /root]#

You will need the IEEE address to configure your storage system.

20. Optionally, check to see that the port is correctly configured on the switch.

The procedure you need to use is different for each switch vendor. We show the procedure for the IBM 2109 family of switches in Example 7-8. For INRANGE switches, you should use the IN-VSN Enterprise Manager application to find the same information. For McDATA switches, you can find the information in the EFC Manager application.

Example 7-8 Checking the Emulex adapter switch connectivity

```
[root@clyde /root]# telnet 9.1.38.159
Trying 9.1.38.159...
Connected to 9.1.38.159.
Escape character is '^'.
```

```
Fabric OS (tm) Release a2.4.1c
```

```
login: admin
Password:
```

```
itsosw3:admin> switchShow
switchName:    itsosw3
switchType:    2.4
switchState:   Online
switchRole:    Principal
switchDomain:  3
switchId:     fffc03
```

```

switchWwn:      10:00:00:60:69:10:64:cf
switchBeacon:   OFF
port 0: id Online      F-Port  21:00:00:e0:8b:02:a7:2d
port 1: -- No_Module
port 2: id Online      E-Port  10:00:00:60:69:20:1d:4e "itsosw1" (downstream)
port 3: -- No_Module
port 4: id No_Light
port 5: id Online      F-Port  10:00:00:00:c9:23:c3:05
port 6: id Online      E-Port  10:00:00:60:69:20:1d:4e "itsosw1"
port 7: id No_Light
port 8: sw Online      F-Port  10:00:00:00:c9:22:bf:75
port 9: -- No_Module
port 10: sw No_Light
port 11: sw No_Light
port 12: sw Online     L-Port  1 public
port 13: id No_Light
port 14: id Online     F-Port  10:00:00:00:c9:25:8e:54
port 15: id No_Light
itsosw3:admin> exit
Connection closed by foreign host.
[root@clyde /root]#

```

From Example 7-8, we see that the port we are using (port 5) is online, it is an F_Port (for fabric attachment), and the WWN listed is the same that we see on our adapter in Example 7-7.

21. Your FC HBA is now correctly installed.

7.5.4 Configuring storage for the Emulex FC HBA

After the adapter installation, we have to create any necessary switch zones and allocate the storage for our host from the storage system we are using. As both of those steps are different for different switch manufacturers and different storage subsystems, they are not described here.

We have access to a single volume in a storage subsystem. We want to configure that volume to our host system. The procedure is as follows:

1. Reboot the machine and log in as root.

Note: We also tried to avoid rebooting, when adding storage, by unloading and reloading the Emulex driver. This consistently caused the Linux kernel to panic.

2. Load the Emulex driver:

```

[root@clyde /root]# modprobe lpfcdd
[root@clyde /root]#

```

3. Check the kernel messages with the `dmesg` command:

```
[root@clyde /root]# dmesg
...
Emulex LightPulse FC SCSI/IP 4.11e
!lpfc-2080374785:045:Vital Product Data Data: 82 23 0 36
!lpfc-2080374785:031:Link Up Event received Data: 1 1 1 0
scsi2 : Emulex LPFC (DRAGONFLY) SCSI on PCI bus 01 device 20 irq 11
Vendor: DEC Model: HSG80 Rev: V85S
Type: Direct-Access ANSI SCSI revision: 02
Attached scsi disk sdb at scsi2, channel 0, id 0, lun 0
SCSI device sdb: 17418235 512-byte hdwr sectors (8918 MB)
sdb:
[root@clyde /root]#
```

Note that now we also see the ID of our disk (sdb). If you already had partitions on the disk, they would also be shown.

4. Start using the new LUN like any normal SCSI disk.

We tested the Fibre Channel drive according to the procedure in Example 7-9:

Example 7-9 Testing a Fibre Channel disk with Emulex adapter

```
[root@clyde /root]# fdisk /dev/sdb
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel. Changes will remain in memory only,
until you decide to write them. After that, of course, the previous
content won't be recoverable.
```

The number of cylinders for this disk is set to 8504.
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): **p**

Disk /dev/sdb: 255 heads, 64 sectors, 1084 cylinders
Units = cylinders of 16065 * 512 bytes

Device	Boot	Start	End	Blocks	Id	System
--------	------	-------	-----	--------	----	--------

Command (m for help): **n**

Command action

e extended
p primary partition (1-4)

p

Partition number (1-4): **1**

First cylinder (1-8504, default 1): **1**

Last cylinder or +size or +sizeM or +sizeK (1-1084, default 1084): **1084**

```

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: If you have created or modified any DOS 6.x
partitions, please see the fdisk manual page for additional
information.
Syncing disks.
[root@clyde /root]# mkfs /dev/sdb1
mke2fs 1.19, 13-Jul-2000 for EXT2 FS 0.5b, 95/08/09
Filesystem label=
OS type: Linux
Fragment size=4096 (log=2)
1089152 inodes, 2177020 blocks
108851 blocks (5.00%) reserved for the super user
First data block=0
67 block groups
32768 blocks per group, 32768 fragments per group
16256 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Writing inode tables: done
Writing superblocks and filesystem accounting information: done
[root@clyde /root]# mkdir /test
[root@clyde /root]# mount /dev/sdb1 /test
[root@clyde /root]# df -k
Filesystem          1k-blocks      Used Available Use% Mounted on
/dev/sda7            402885        85376   296708   23% /
/dev/sda2             201433         3822   187210    3% /boot
/dev/sda5            3024016       540128  2330272   19% /usr
/dev/sdb1            8570444         20   8135068    1% /test
[root@clyde /root]# cp /etc/hosts /test
[root@clyde /root]# ls -l /etc/hosts /test/hosts
-rw-r--r--  1 root   root           181 Dec  3 04:00 /etc/hosts
-rw-r--r--  1 root   root           181 Dec  3 15:15 /test/hosts
[root@clyde /root]# diff /etc/hosts /test/hosts
[root@clyde /root]#

```

The connection is successful.

Glossary

8B/10B A data encoding scheme developed by IBM, translating byte-wide data to an encoded 10-bit format. Fibre Channel's FC-1 level defines this as the method to be used to encode and decode data transmissions over the Fibre channel.

Adapter A hardware unit that aggregates other I/O units, devices or communications links to a system bus.

ADSM ADSTAR Distributed Storage Manager.

Agent (1) In the client-server model, the part of the system that performs information preparation and exchange on behalf of a client or server application. (2) In SNMP, the word agent refers to the managed system. See also: Management Agent

AIT Advanced Intelligent Tape - A magnetic tape format by Sony that uses 8mm cassettes, but is only used in specific drives.

AL See Arbitrated Loop

ANSI American National Standards Institute - The primary organization for fostering the development of technology standards in the United States. The ANSI family of Fibre Channel documents provide the standards basis for the Fibre Channel architecture and technology. See FC-PH

Arbitration The process of selecting one respondent from a collection of several candidates that request service concurrently.

Arbitrated Loop A Fibre Channel interconnection technology that allows up to 126 participating node ports and one participating fabric port to communicate.

ATL Automated Tape Library - Large scale tape storage system, which uses multiple tape drives and mechanisms to address 50 or more cassettes.

ATM Asynchronous Transfer Mode - A type of packet switching that transmits fixed-length units of data.

Backup A copy of computer data that is used to recreate data that has been lost, mislaid, corrupted, or erased. The act of creating a copy of computer data that can be used to recreate data that has been lost, mislaid, corrupted or erased.

Bandwidth Measure of the information capacity of a transmission channel.

Bridge (1) A component used to attach more than one I/O unit to a port. (2) A data communications device that connects two or more networks and forwards packets between them. The bridge may use similar or dissimilar media and signaling systems. It operates at the data link level of the OSI model. Bridges read and filter data packets and frames.

Bridge/Router A device that can provide the functions of a bridge, router or both concurrently. A bridge/router can route one or more protocols, such as TCP/IP, and bridge all other traffic. See also: Bridge, Router

Broadcast Sending a transmission to all N_Ports on a fabric.

Channel A point-to-point link, the main task of which is to transport data from one point to another.

Channel I/O A form of I/O where request and response correlation is maintained through some form of source, destination and request identification.

CIFS Common Internet File System

Class of Service A Fibre Channel frame delivery scheme exhibiting a specified set of delivery characteristics and attributes.

Class-1 A class of service providing dedicated connection between two ports with confirmed delivery or notification of non-deliverability.

Class-2 A class of service providing a frame switching service between two ports with confirmed delivery or notification of non-deliverability.

Class-3 A class of service providing frame switching datagram service between two ports or a multicast service between a multicast originator and one or more multicast recipients.

Class-4 A class of service providing a fractional bandwidth virtual circuit between two ports with confirmed delivery or notification of non-deliverability.

Class-6 A class of service providing a multicast connection between a multicast originator and one or more multicast recipients with confirmed delivery or notification of non-deliverability.

Client A software program used to contact and obtain data from a *server* software program on another computer - often across a great distance. Each *client* program is designed to work specifically with one or more kinds of server programs and each server requires a specific kind of client program.

Client/Server The relationship between machines in a communications network. The client is the requesting machine, the server the supplying machine. Also used to describe the information management relationship between software components in a processing system.

Cluster A type of parallel or distributed system that consists of a collection of interconnected whole computers and is used as a single, unified computing resource.

Coaxial Cable A transmission media (cable) used for high speed transmission. It is called *coaxial* because it includes one physical channel that carries the signal surrounded (after a layer of insulation) by another concentric physical channel, both of which run along the same axis. The inner channel carries the signal and the outer channel serves as a ground.

Controller A component that attaches to the system topology through a channel semantic protocol that includes some form of request/response identification.

CRC Cyclic Redundancy Check - An error-correcting code used in Fibre Channel.

DASD Direct Access Storage Device - any on-line storage device: a disc, drive or CD-ROM.

DAT Digital Audio Tape - A tape media technology designed for very high quality audio recording and data backup. DAT cartridges look like audio cassettes and are often used in mechanical auto-loaders. typically, a DAT cartridge provides 2GB of storage. But new DAT systems have much larger capacities.

Data Sharing A SAN solution in which files on a storage device are shared between multiple hosts.

Datagram Refers to the Class 3 Fibre Channel Service that allows data to be sent rapidly to multiple devices attached to the fabric, with no confirmation of delivery.

dB Decibel - a ratio measurement distinguishing the percentage of signal attenuation between the input and output power. Attenuation (loss) is expressed as dB/km

Disk Mirroring A fault-tolerant technique that writes data simultaneously to two hard disks using the same hard disk controller.

Disk Pooling A SAN solution in which disk storage resources are pooled across multiple hosts rather than be dedicated to a specific host.

DLT Digital Linear Tape - A magnetic tape technology originally developed by Digital Equipment Corporation (DEC) and now sold by Quantum. DLT cartridges provide storage capacities from 10 to 35GB.

E_Port Expansion Port - a port on a switch used to link multiple switches together into a Fibre Channel switch fabric.

ECL Emitter Coupled Logic - The type of transmitter used to drive copper media such as Twinax, Shielded Twisted Pair, or Coax.

Enterprise Network A geographically dispersed network under the auspices of one organization.

Entity In general, a real or existing thing from the Latin ens, or being, which makes the distinction between a thing's existence and its qualities. In programming, engineering and probably many other contexts, the word is used to identify units, whether concrete things or abstract ideas, that have no ready name or label.

ESCON Enterprise System Connection

Exchange A group of sequences which share a unique identifier. All sequences within a given exchange use the same protocol. Frames from multiple sequences can be multiplexed to prevent a single exchange from consuming all the bandwidth. See also: Sequence

F_Node Fabric Node - a fabric attached node.

F_Port Fabric Port - a port used to attach a Node Port (N_Port) to a switch fabric.

Fabric Fibre Channel employs a fabric to connect devices. A fabric can be as simple as a single cable connecting two devices. The term is most often used to describe a more complex network utilizing hubs, switches and gateways.

Fabric Login Fabric Login (FLOGI) is used by an N_Port to determine if a fabric is present and, if so, to initiate a session with the fabric by exchanging service parameters with the fabric. Fabric Login is performed by an N_Port following link initialization and before communication with other N_Ports is attempted.

FC Fibre Channel

FC-0 Lowest level of the Fibre Channel Physical standard, covering the physical characteristics of the interface and media

FC-1 Middle level of the Fibre Channel Physical standard, defining the 8B/10B encoding/decoding and transmission protocol.

FC-2 Highest level of the Fibre Channel Physical standard, defining the rules for signaling protocol and describing transfer of frame, sequence and exchanges.

FC-3 The hierarchical level in the Fibre Channel standard that provides common services such as striping definition.

FC-4 The hierarchical level in the Fibre Channel standard that specifies the mapping of upper-layer protocols to levels below.

FCA Fiber Channel Association.

FC-AL Fibre Channel Arbitrated Loop - A reference to the Fibre Channel Arbitrated Loop standard, a shared gigabit media for up to 127 nodes, one of which may be attached to a switch fabric. See also: Arbitrated Loop.

FC-CT Fibre Channel common transport protocol

FC-FG Fibre Channel Fabric Generic - A reference to the document (ANSI X3.289-1996) which defines the concepts, behavior and characteristics of the Fibre Channel Fabric along with suggested partitioning of the 24-bit address space to facilitate the routing of frames.

FC-FP Fibre Channel HIPPI Framing Protocol - A reference to the document (ANSI X3.254-1994) defining how the HIPPI framing protocol is transported via the fibre channel

FC-GS Fibre Channel Generic Services - A reference to the document (ANSI X3.289-1996) describing a common transport protocol used to communicate with the server functions, a full X500 based directory service, mapping of the Simple Network Management Protocol (SNMP) directly to the Fibre Channel, a time server and an alias server.

FC-LE Fibre Channel Link Encapsulation - A reference to the document (ANSI X3.287-1996) which defines how IEEE 802.2 Logical Link Control (LLC) information is transported via the Fibre Channel.

FC-PH A reference to the Fibre Channel Physical and Signaling standard ANSI X3.230, containing the definition of the three lower levels (FC-0, FC-1, and FC-2) of the Fibre Channel.

FC-PLDA Fibre Channel Private Loop Direct Attach - See PLDA.

FC-SB Fibre Channel Single Byte Command Code Set - A reference to the document (ANSI

X.271-1996) which defines how the ESCON command set protocol is transported using the fibre channel.

FC-SW Fibre Channel Switch Fabric - A reference to the ANSI standard under development that further defines the fabric behavior described in FC-FG and defines the communications between different fabric elements required for those elements to coordinate their operations and management address assignment.

FC Storage Director See SAN Storage Director

FCA Fibre Channel Association - a Fibre Channel industry association that works to promote awareness and understanding of the Fibre Channel technology and its application and provides a means for implementers to support the standards committee activities.

FCLC Fibre Channel Loop Association - an independent working group of the Fibre Channel Association focused on the marketing aspects of the Fibre Channel Loop technology.

FCP Fibre Channel Protocol - the mapping of SCSI-3 operations to Fibre Channel.

Fiber Optic Refers to the medium and the technology associated with the transmission of information along a glass or plastic wire or fiber.

Fibre Channel A technology for transmitting data between computer devices at a data rate of up to 4 Gb/s. It is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

FICON Fibre Connection - A next-generation I/O solution for IBM S/390 parallel enterprise server.

FL_Port Fabric Loop Port - the access point of the fabric for physically connecting the user's Node Loop Port (NL_Port).

FLOGI See Fabric Log In

Frame A linear set of transmitted bits that define the basic transport unit. The frame is the most basic element of a message in Fibre Channel communications, consisting of a 24-byte header and zero to 2112 bytes of data. See also: Sequence

FSP Fibre Channel Service Protocol - The common FC-4 level protocol for all services, transparent to the fabric type or topology.

Full-Duplex A mode of communications allowing simultaneous transmission and reception of frames.

G_Port Generic Port - a generic switch port that is either a Fabric Port (F_Port) or an Expansion Port (E_Port). The function is automatically determined during login.

Gateway A node on a network that interconnects two otherwise incompatible networks.

Gb/s Gigabits per second. Also sometimes referred to as Gbps. In computing terms it is approximately 1,000,000,000 bits per second. Most precisely it is 1,073,741,824 (1024 x 1024 x 1024) bits per second.

GB/s Gigabytes per second. Also sometimes referred to as GBps. In computing terms it is approximately 1,000,000,000 bytes per second. Most precisely it is 1,073,741,824 (1024 x 1024 x 1024) bytes per second.

GBIC GigaBit Interface Converter - Industry standard transceivers for connection of Fibre Channel nodes to arbitrated loop hubs and fabric switches.

Gigabit One billion bits, or one thousand megabits.

GLM Gigabit Link Module - a generic Fibre Channel transceiver unit that integrates the key functions necessary for installation of a Fibre channel media interface on most systems.

Half-Duplex A mode of communications allowing either transmission or reception of frames at any point in time, but not both (other than link control frames which are always permitted).

Hardware The mechanical, magnetic and electronic components of a system, e.g., computers, telephone switches, terminals and the like.

HBA Host Bus Adapter

HIPPI High Performance Parallel Interface - An ANSI standard defining a channel that transfers

data between CPUs and from a CPU to disk arrays and other peripherals.

HMMP HyperMedia Management Protocol

HMMS HyperMedia Management Schema - the definition of an implementation-independent, extensible, common data description/schema allowing data from a variety of sources to be described and accessed in real time regardless of the source of the data. See also: WEBM, HMMP

HSM Hierarchical Storage Management - A software and hardware system that moves files from disk to slower, less expensive storage media based on rules and observation of file activity. Modern HSM systems move files from magnetic disk to optical disk to magnetic tape.

HUB A Fibre Channel device that connects nodes into a logical loop by using a physical star topology. Hubs will automatically recognize an active node and insert the node into the loop. A node that fails or is powered off is automatically removed from the loop.

HUB Topology see Loop Topology

Hunt Group A set of associated Node Ports (N_Ports) attached to a single node, assigned a special identifier that allows any frames containing this identifier to be routed to any available Node Port (N_Port) in the set.

In-Band Signaling This is signaling that is carried in the same channel as the information. Also referred to as inband.

Information Unit A unit of information defined by an FC-4 mapping. Information Units are transferred as a Fibre Channel Sequence.

Intermix A mode of service defined by Fibre Channel that reserves the full Fibre Channel bandwidth for a dedicated Class 1 connection, but also allows connection-less Class 2 traffic to share the link if the bandwidth is available.

I/O Input/output

IP Internet Protocol

IPI Intelligent Peripheral Interface

Isochronous Transmission Data transmission which supports network-wide timing

requirements. A typical application for isochronous transmission is a broadcast environment which needs information to be delivered at a predictable time.

JBOD Just a bunch of disks.

Jukebox A device that holds multiple optical disks and one or more disk drives, and can swap disks in and out of the drive as needed.

L_Port Loop Port - A node or fabric port capable of performing Arbitrated Loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.

LAN See Local Area Network - A network covering a relatively small geographic area (usually not larger than a floor or small building). Transmissions within a Local Area Network are mostly digital, carrying data among stations at rates usually above one megabit/s.

Latency A measurement of the time it takes to send a frame between two locations.

LC Lucent Connector. A registered trademark of Lucent Technologies.

Link A connection between two Fibre Channel ports consisting of a transmit fibre and a receive fibre.

Link_Control_Facility A termination card that handles the logical and physical control of the Fibre Channel link for each mode of use.

LIP A Loop Initialization Primitive sequence is a special fibre channel sequence that is used to start loop initialization. Allows ports to establish their port addresses.

Local Area Network (LAN) A network covering a relatively small geographic area (usually not larger than a floor or small building). Transmissions within a Local Area Network are mostly digital, carrying data among stations at rates usually above one megabit/s.

Login Server Entity within the Fibre Channel fabric that receives and responds to login requests.

Loop Circuit A temporary point-to-point like path that allows bi-directional communications between loop-capable ports.

Loop Topology An interconnection structure in which each point has physical links to two neighbors resulting in a closed circuit. In a loop topology, the available bandwidth is shared.

LVD Low Voltage Differential

Management Agent A process that exchanges a managed node's information with a management station.

Managed Node A managed node is a computer, a storage system, a gateway, a media device such as a switch or hub, a control instrument, a software product such as an operating system or an accounting package, or a machine on a factory floor, such as a robot.

Managed Object A variable of a managed node. This variable contains one piece of information about the node. Each node can have several objects.

Management Station A host system that runs the management software.

Mb/s Megabits per second. Also sometimes referred to as Mbps. In computing terms it is approximately 1,000,000 bits per second. Most precisely it is 1,048,576 (1024 x 1024) bits per second.

MB/s Megabytes per second. Also sometimes referred to as MBps. In computing terms it is approximately 1,000,000 bytes per second. Most precisely it is 1,048,576 (1024 x 1024) bytes per second.

Meter 39.37 inches, or just slightly larger than a yard (36 inches)

Media Plural of medium. The physical environment through which transmission signals pass. Common media include copper and fiber optic cable.

Media Access Rules (MAR).

MIA Media Interface Adapter - MIAs enable optic-based adapters to interface to

copper-based devices, including adapters, hubs, and switches.

MIB Management Information Block - A formal description of a set of network objects that can be managed using the Simple Network Management Protocol (SNMP). The format of the MIB is defined as part of SNMP and is a hierarchical structure of information relevant to a specific device, defined in object oriented terminology as a collection of objects, relations, and operations among objects.

Mirroring The process of writing data to two separate physical devices simultaneously.

MM Multi-Mode - See Multi-Mode Fiber

MMF See Multi-Mode Fiber - - In optical fiber technology, an optical fiber that is designed to carry multiple light rays or modes concurrently, each at a slightly different reflection angle within the optical core. Multi-Mode fiber transmission is used for relatively short distances because the modes tend to disperse over longer distances. See also: Single-Mode Fiber, SMF

Multicast Sending a copy of the same transmission from a single source device to multiple destination devices on a fabric. This includes sending to all N_Ports on a fabric (broadcast) or to only a subset of the N_Ports on a fabric (multicast).

Multi-Mode Fiber (MMF) In optical fiber technology, an optical fiber that is designed to carry multiple light rays or modes concurrently, each at a slightly different reflection angle within the optical core. Multi-Mode fiber transmission is used for relatively short distances because the modes tend to disperse over longer distances. See also: Single-Mode Fiber

Multiplex The ability to intersperse data from multiple sources and destinations onto a single transmission medium. Refers to delivering a single transmission to multiple destination Node Ports (N_Ports).

N_Port Node Port - A Fibre Channel-defined hardware entity at the end of a link which provides the mechanisms necessary to transport information units to or from another node.

N_Port Login N_Port Login (PLOGI) allows two N_Ports to establish a session and exchange identities and service parameters. It is performed following completion of the fabric login process and prior to the FC-4 level operations with the destination port. N_Port Login may be either explicit or implicit.

Name Server Provides translation from a given node name to one or more associated N_Port identifiers.

NAS Network Attached Storage - a term used to describe a technology where an integrated storage system is attached to a messaging network that uses common communications protocols, such as TCP/IP.

NDMP Network Data Management Protocol

Network An aggregation of interconnected nodes, workstations, file servers, and/or peripherals, with its own protocol that supports interaction.

Network Topology Physical arrangement of nodes and interconnecting communications links in networks based on application requirements and geographical distribution of users.

NFS Network File System - A distributed file system in UNIX developed by Sun Microsystems which allows a set of computers to cooperatively access each other's files in a transparent manner.

NL_Port Node Loop Port - a node port that supports Arbitrated Loop devices.

NMS Network Management System - A system responsible for managing at least part of a network. NMSs communicate with agents to help keep track of network statistics and resources.

Node An entity with one or more N_Ports or NL_Ports.

Non-Blocking A term used to indicate that the capabilities of a switch are such that the total number of available transmission paths is equal to the number of ports. Therefore, all ports can have simultaneous access through the switch.

Non-L_Port A Node or Fabric port that is not capable of performing the Arbitrated Loop

functions and protocols. N_Ports and F_Ports are not loop-capable ports.

Operation A term defined in FC-2 that refers to one of the Fibre Channel *building blocks* composed of one or more, possibly concurrent, exchanges.

Optical Disk A storage device that is written and read by laser light.

Optical Fiber A medium and the technology associated with the transmission of information as light pulses along a glass or plastic wire or fiber.

Ordered Set A Fibre Channel term referring to four 10-bit characters (a combination of data and special characters) providing low-level link functions, such as frame demarcation and signaling between two ends of a link.

Originator A Fibre Channel term referring to the initiating device.

Out of Band Signaling This is signaling that is separated from the channel carrying the information. Also referred to as outband.

Peripheral Any computer device that is not part of the essential computer (the processor, memory and data paths) but is situated relatively close by. A near synonym is input/output (I/O) device.

Petard A device that is small and sometimes explosive.

PLDA Private Loop Direct Attach - A technical report which defines a subset of the relevant standards suitable for the operation of peripheral devices such as disks and tapes on a private loop.

PLOGI See N_Port Login

Point-to-Point Topology An interconnection structure in which each point has physical links to only one neighbor resulting in a closed circuit. In point-to-point topology, the available bandwidth is dedicated.

Port The hardware entity within a node that performs data communications over the Fibre Channel.

Port Bypass Circuit A circuit used in hubs and disk enclosures to automatically open or close the loop to add or remove nodes on the loop.

Private NL_Port An NL_Port which does not attempt login with the fabric and only communicates with other NL Ports on the same loop.

Protocol A data transmission convention encompassing timing, control, formatting and data representation.

Public NL_Port An NL_Port that attempts login with the fabric and can observe the rules of either public or private loop behavior. A public NL_Port may communicate with both private and public NL_Ports.

Quality of Service (QoS) A set of communications characteristics required by an application. Each QoS defines a specific transmission priority, level of route reliability, and security level.

RAID Redundant Array of Inexpensive or Independent Disks. A method of configuring multiple disk drives in a storage subsystem for high availability and high performance.

Raid 0 Level 0 RAID support - Striping, no redundancy

Raid 1 Level 1 RAID support - mirroring, complete redundancy

Raid 5 Level 5 RAID support, Striping with parity

Repeater A device that receives a signal on an electromagnetic or optical transmission medium, amplifies the signal, and then retransmits it along the next leg of the medium.

Responder A Fibre Channel term referring to the answering device.

Router (1) A device that can decide which of several paths network traffic will follow based on some optimal metric. Routers forward packets from one network to another based on network-layer information. (2) A dedicated computer hardware and/or software package which manages the connection between two or more networks. See also: Bridge, Bridge/Router

SAF-TE SCSI Accessed Fault-Tolerant Enclosures

SAN A Storage Area Network (SAN) is a dedicated, centrally managed, secure information infrastructure, which enables any-to-any interconnection of servers and storage systems.

SAN System Area Network - term originally used to describe a particular symmetric multiprocessing (SMP) architecture in which a switched interconnect is used in place of a shared bus. Server Area Network - refers to a switched interconnect between multiple SMPs.

SC Connector A fiber optic connector standardized by ANSI TIA/EIA-568A for use in structured wiring installations.

Scalability The ability of a computer application or product (hardware or software) to continue to function well as it (or its context) is changed in size or volume. For example, the ability to retain performance levels when adding additional processors, memory and/or storage.

SCSI Small Computer System Interface - A set of evolving ANSI standard electronic interfaces that allow personal computers to communicate with peripheral hardware such as disk drives, tape drives, CD_ROM drives, printers and scanners faster and more flexibly than previous interfaces. The table below identifies the major characteristics of the different SCSI version.

SCSI Version	Signal Rate MHz	Bus-Width (bits)	Max. DTR (MB/s)	Max. Num. Devices	Max. Cable Length (m)
SCSI -1	5	8	5	7	6
SCSI -2	5	8	5	7	6
Wide SCSI -2	5	16	10	15	6
Fast SCSI -2	10	8	10	7	6
Fast Wide SCSI -2	10	16	20	15	6

Ultra SCSI	20	8	20	7	1.5
Ultra SCSI -2	20	16	40	7	12
Ultra 2 LVD SCSI	40	16	80	15	12

SCSI-3 SCSI-3 consists of a set of primary commands and additional specialized command sets to meet the needs of specific device types. The SCSI-3 command sets are used not only for the SCSI-3 parallel interface but for additional parallel and serial protocols, including Fibre Channel, Serial Bus Protocol (used with IEEE 1394 Firewire physical protocol) and the Serial Storage Protocol (SSP).

SCSI-FCP The term used to refer to the ANSI Fibre Channel Protocol for SCSI document (X3.269-199x) that describes the FC-4 protocol mappings and the definition of how the SCSI protocol and command set are transported using a Fibre Channel interface.

Sequence A series of frames strung together in numbered order which can be transmitted over a Fibre Channel connection as a single operation. See also: Exchange

SERDES Serializer Deserializer

Server A computer which is dedicated to one task.

SES SCSI Enclosure Services - ANSI SCSI-3 proposal that defines a command set for soliciting basic device status (temperature, fan speed, power supply status, etc.) from a storage enclosures.

Single-Mode Fiber In optical fiber technology, an optical fiber that is designed for the transmission of a single ray or mode of light as a carrier. It is a single light path used for long-distance signal transmission. See also: Multi-Mode Fiber

SMART Self Monitoring and Reporting Technology

SM Single Mode - See Single-Mode Fiber

SMF Single-Mode Fiber - In optical fiber technology, an optical fiber that is designed for the transmission of a single ray or mode of light as a carrier. It is a single light path used for long-distance signal transmission. See also: MMF

SNIA Storage Networking Industry Association. A non-profit organization comprised of more than 77 companies and individuals in the storage industry.

SN Storage Network. See also: SAN

SNMP Simple Network Management Protocol - The Internet network management protocol which provides a means to monitor and set network configuration and run-time parameters.

SNMWG Storage Network Management Working Group is chartered to identify, define and support open standards needed to address the increased management requirements imposed by storage area network environments.

SSA Serial Storage Architecture - A high speed serial loop-based interface developed as a high speed point-to-point connection for peripherals, particularly high speed storage arrays, RAID and CD-ROM storage by IBM.

Star The physical configuration used with hubs in which each user is connected by communications links radiating out of a central hub that handles all communications.

StorWatch Expert These are StorWatch applications that employ a 3 tiered architecture that includes a management interface, a StorWatch manager and agents that run on the storage resource(s) being managed. Expert products employ a StorWatch data base that can be used for saving key management data (e.g. capacity or performance metrics). Expert products use the agents as well as analysis of storage data saved in the data base to perform higher value functions including - reporting of capacity, performance, etc. over time (trends), configuration of multiple devices based on policies, monitoring of capacity and performance, automated responses to events or conditions, and storage related data mining.

StorWatch Specialist A StorWatch interface for managing an individual fibre Channel device or a limited number of like devices (that can be viewed as a single group). StorWatch specialists typically provide simple, point-in-time management functions such as configuration, reporting on asset and status information, simple device and event monitoring, and perhaps some service utilities.

Striping A method for achieving higher bandwidth using multiple N_Ports in parallel to transmit a single information unit across multiple levels.

STP Shielded Twisted Pair

Storage Media The physical device itself, onto which data is recorded. Magnetic tape, optical disks, floppy disks are all storage media.

Switch A component with multiple entry/exit points (ports) that provides dynamic connection between any two of these points.

Switch Topology An interconnection structure in which any entry point can be dynamically connected to any exit point. In a switch topology, the available bandwidth is scalable.

T11 A technical committee of the National Committee for Information Technology Standards, titled T11 I/O Interfaces. It is tasked with developing standards for moving data in and out of computers.

Tape Backup Making magnetic tape copies of hard disk and optical disc files for disaster recovery.

Tape Pooling A SAN solution in which tape resources are pooled and shared across multiple hosts rather than being dedicated to a specific host.

TCP Transmission Control Protocol - a reliable, full duplex, connection-oriented end-to-end transport protocol running on top of IP.

TCP/IP Transmission Control Protocol/ Internet Protocol - a set of communications protocols that support peer-to-peer connectivity functions for both local and wide area networks.

Time Server A Fibre Channel-defined service function that allows for the management of all timers used within a Fibre Channel system.

Topology An interconnection scheme that allows multiple Fibre Channel ports to communicate. For example, point-to-point, Arbitrated Loop, and switched fabric are all Fibre Channel topologies.

T_Port An ISL port more commonly known as an E_Port, referred to as a Trunk port and used by INRANGE.

TL_Port A private to public bridging of switches or directors, referred to as Translative Loop.

Twinax A transmission media (cable) consisting of two insulated central conducting leads of coaxial cable.

Twisted Pair A transmission media (cable) consisting of two insulated copper wires twisted around each other to reduce the induction (thus interference) from one wire to another. The twists, or lays, are varied in length to reduce the potential for signal interference between pairs. Several sets of twisted pair wires may be enclosed in a single cable. This is the most common type of transmission media.

ULP Upper Level Protocols

UTC Under-The-Covers, a term used to characterize a subsystem in which a small number of hard drives are mounted inside a higher function unit. The power and cooling are obtained from the system unit. Connection is by parallel copper ribbon cable or pluggable backplane, using IDE or SCSI protocols.

UTP Unshielded Twisted Pair

Virtual Circuit A unidirectional path between two communicating N_Ports that permits fractional bandwidth.

WAN Wide Area Network - A network which encompasses inter-connectivity between devices over a wide geographic area. A wide area network may be privately owned or rented, but the term usually connotes the inclusion of public (shared) networks.

WDM Wave Division Multiplexing - A technology that puts data from different sources together on an optical fiber, with each signal carried on its own separate light wavelength. Using WDM, up to 80 (and theoretically more) separate wavelengths or channels of data can be multiplexed into a stream of light transmitted on a single optical fiber.

WEBM Web-Based Enterprise Management - A consortium working on the development of a series of standards to enable active management and monitoring of network-based elements.

Zoning In Fibre Channel environments, the grouping together of multiple ports to form a virtual private storage network. Ports that are members of a group or zone can communicate with each other but are isolated from ports in other zones.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 180.

- ▶ *IBM SAN Survival Guide*, SG24-6143
- ▶ *IBM SAN Survival Guide Featuring the IBM 2109*, SG24-6127
- ▶ *IBM SAN Survival Guide Featuring the McDATA Portfolio*, SG24-6149
- ▶ *IBM SAN Survival Guide Featuring the INRANGE Portfolio*, SG24-6150
- ▶ *Designing an IBM Storage Area Network*, SG24-5758
- ▶ *Implementing an Open IBM SAN*, SG24-6116
- ▶ *Introduction to Storage Area Network, SAN*, SG24-5470
- ▶ *IBM Storage Solutions for Server Consolidation*, SG24-5355
- ▶ *Implementing the Enterprise Storage Server in Your Environment*, SG24-5420
- ▶ *Implementing Linux with IBM Disk Storage*, SG24-6261
- ▶ *Storage Area Networks: Tape Future In Fabrics*, SG24-5474
- ▶ *IBM Enterprise Storage Server*, SG24-5465

Other resources

These publications are also relevant as further information sources:

- ▶ *Building Storage Networks*, ISBN 0072120509

These IBM publications are also relevant as further information sources:

- ▶ *ESS Web Interface User's Guide for ESS Specialist and ESS Copy Services*, SC26-7346
- ▶ *IBM Storage Area Network Data Gateway Installation and User's Guide*, SC26-7304

- ▶ *IBM Enterprise Storage Server Configuration Planner, SC26-7353*
- ▶ *IBM Enterprise Storage Server Quick Configuration Guide, SC26-7354*
- ▶ *IBM SAN Fibre Channel Managed Hub 3534 Service Guide, SY27-7616*
- ▶ *IBM Enterprise Storage Server Introduction and Planning Guide, 2105 Models E10, E20, F10 and F20, GC26-7294*
- ▶ *IBM Enterprise Storage Server User's Guide, 2105 Models E10, E20, F10 and F20, SC26-7295*
- ▶ *IBM Enterprise Storage Server Host Systems Attachment Guide, 2105 Models E10, E20, F10 and F20, SC26-7296*
- ▶ *IBM Enterprise Storage Server SCSI Command Reference, 2105 Models E10, E20, F10 and F20, SC26-7297*
- ▶ *IBM Enterprise Storage Server System/390 Command Reference, 2105 Models E10, E20, F10 and F20, SC26-7298*
- ▶ *IBM Storage Solutions Safety Notices, GC26-7229*
- ▶ *PCI Adapter Placement Reference, SA38-0583*
- ▶ *Translated External Devices/Safety Information, SA26-7003*
- ▶ *Electrical Safety for IBM Customer Engineers, S229-8124*
- ▶ *SLIC Router Installation and Users Guide, 310-605759*
- ▶ *SLIC Manager Installation and User Guide, 310-605807*

The IBM NUMA-Q publications which are also relevant as further information sources are available on the Web at:

<http://webdocs.numaq.ibm.com>

The JNI publications which are also relevant as further information sources are available on the Web at:

<http://www.jni.com/Support/installguides.cfm>

These McDATA publications are also relevant as further information sources:

- ▶ *ED-5000 Director Planning Manual, 620-005000*
- ▶ *Enterprise Fabric Connectivity Manager User Manual, 620-005001*
- ▶ *ED-5000 Director User Manual, 620-005002*
- ▶ *ED-5000 Director Service Manual, 620-005004*
- ▶ *ED-6064 Director Planning Manual, 620-000106-100*
- ▶ *ED-6064 Director User Manual, 620-000107*
- ▶ *ED-6064 Director Installation and Service Manual, 620-000108*

- ▶ *Enterprise Fabric Connectivity Manager User Manual*, 620-005001
- ▶ *FC-512 Fabriccenter Equipment Cabinet Installation and Service Manual*, 620-000100
- ▶ *ES-3016 Switch Planning Manual*, 620-000110-100
- ▶ *ES-3016 Switch User Manual*, 620-000111
- ▶ *ES-3016 Switch Installation and Service Manual*, 620-000112
- ▶ *ES-3032 Switch Planning Manual*, 620-000118-000
- ▶ *ES-3032 Switch User Manual*, 620-000117-000
- ▶ *ES-3032 Switch Installation and Service Manual*, 620-000116-000
- ▶ *ES-1000 Switch Planning Manual*, 620-000102-000
- ▶ *ES-1000 Switch User Manual*, 620-000103
- ▶ *ES-1000 Switch Installation and Service Manual*, 620-000105

These QLogic publications are also relevant as further information sources:

- ▶ *QLA2200 Hardware Manual*, FC0151103-00
- ▶ *QLA2200 Hardware Manual*, FC0151103-00
- ▶ *QLA2100 Software Manual*, FC0153301-00
- ▶ *QLA2100 Hardware Manual*, FC0151102-00
- ▶ *QMS V1 Installation Guide*, FC0051104-00
- ▶ *QLview for Fibre Operations Guide*, FC0051101-00
- ▶ *QLconfig Operations Guide*, FC0051102-00

These Compaq publications are also relevant as further information sources:

- ▶ *64-Bit PCI to Fibre Channel Host Bus Adapter Release Notes for Tru64 UNIX and OpenVMS*, AV-RLLUA-TE
- ▶ *64-Bit PCI to Fibre Channel Host Bus Adapter User Guide*, AA-RKPDB-TE

Referenced Web sites

These Web sites are also relevant as further information sources:

- ▶ <http://www.storage.ibm.com/ibmsan/index.htm>
IBM Enterprise SAN
- ▶ <http://www.pc.ibm.com/ww/netfinity/san>
IBM Storage Area Networks: Nefinity Servers
- ▶ <http://www.storage.ibm.com/hardsoft/products/sangateway/supserver.htm>
IBM SAN Data Gateway
- ▶ <http://www.storage.ibm.com/hardsoft/products/tape/ro3superserver.htm>
IBM SAN Data Gateway Router
- ▶ <http://www.storage.ibm.com/hardsoft/products/fcss/fcss.htm>
IBM Fibre Channel RAID Storage Server
- ▶ <http://www.storage.ibm.com/hardsoft/products/ess/ess.htm>
Enterprise Storage Server
- ▶ http://www-1.ibm.com/servers/eserver/pseries/library/hardware_docs/index.html
IBM eServer pSeries and RS/6000 Hardware Documentation
- ▶ <http://www.techsupport.services.ibm.com/server/fixes?view=pSeries>
IBM eServer pSeries and RS/6000 Support
- ▶ <http://www.techsolutions.hp.com>
Hewlett-Packard Support
- ▶ <http://www.software.hp.com>
Hewlett-Packard Software Depot
- ▶ <http://www.redhat.com>
Red Hat, Inc.
- ▶ <http://www.suse.com>
SuSE Inc.
- ▶ <http://www.caldera.com>
Caldera International, Inc.
- ▶ <http://www.turbolinux.com>
Turbolinux Inc.
- ▶ <http://www.cdp.com>
Columbia Data Products, Inc.
- ▶ <http://www.emulex.com>
Emulex Corporation
- ▶ <http://www.fibrechannel.com>
Fibre Channel Industry Association

- ▶ <http://www.jni.com>
JNI Corporation
- ▶ <http://www.mcdata.com>
McDATA Corporation
- ▶ <http://www.pathlight.com>
Pathlight
- ▶ <http://www.qlogic.com>
QLogic Corporation
- ▶ <http://www.sanergy.com>
Tivoli SANergy
- ▶ <http://www.snia.org>
Storage Networking Industry Association
- ▶ <http://www.tivoli.com>
Tivoli
- ▶ <http://www.t11.org>
Technical Committee T11
- ▶ <http://www.vicom.com>
Vicom Systems
- ▶ <http://www.vixel.com>
Vixel
- ▶ <http://www.scsita.org>
SCSI Trade Association
- ▶ <http://www.futureio.org>
InfiniBand (SM) Trade Association
- ▶ <http://www.nsic.org>
National Storage Industry Consortium
- ▶ <http://www.ietf.org>
Internet Engineering Task Force
- ▶ <http://www.ansi.org>
American National Standards Institute
- ▶ <http://www.standards.ieee.org>
Institute of Electrical and Electronics Engineers
- ▶ <http://www.pc.ibm.com/us>
US Personal Systems Group
- ▶ <http://www.compaq.com/products/storageworks/adapters.html>
Compaq Storage - Host Bus Adapters

How to get IBM Redbooks

Search for additional Redbooks or redpieces, view, download, or order hardcopy from the Redbooks Web site:

ibm.com/redbooks

Also download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

Redpieces are Redbooks in progress; not all Redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

Numerics

2031 8, 98, 145
2031-016 9, 99, 145
2031-032 9, 99, 145
2031-L00 9, 99, 145–146
2032 8, 98, 145
2032-001 99, 145
2032-064 99, 145
2032-C36 98, 145
2042 97, 144
2042-C40 98, 144
2105 Enterprise Storage Server 114–115
2108 114
2109 6–7, 38, 69, 96–97, 102, 131, 142, 144, 146,
150, 159
2109-S08 97, 129
2109-S16 129
3590E 115
6227 35, 41, 44
6228 35, 41, 44
64-Bit PCI 128
7556 109
7597 109
7598 109
7599 109

A

A5158A 94
A6684A 94
A6685A 94
Access Restricted 122
adapter 11
adapter configuration 135
adapter configuration settings 16
adapter status 22
AIX 34
alarm notification 13
AlphaServer 127, 129, 131–132, 134, 136–137
Alternate Pathing 61
Arbitrated Loop 108
Arbitrated-Loop 5, 129
attenuation 7, 76, 96, 116, 130, 143
availability 125

B

bandwidth 129
BIOS 15
boot 16, 118
boot probe 118
brand loyalty 2
bridge port 9, 39, 99, 146
bridging 172

C

cable lengths 116
Caldera 140
cascaded 74, 108, 114
cascading 97
CFWDL FW Bundle 118
class 2 108
class 3 108
clean 96
cluster device 4
Compaq 64-bit PCI 128
Compaq AlphaServer 127–128
Compaq KGPSA 128
configShow 78
configuration files 136
configuration information 13
connections 96, 130
connectivity check 136
connectors 7, 130
correct software 117

D

device adapters 6
diagnostics 13
dirt 122
Disable Translative Mode 126
disk storage systems 146
domain 126
domain ids 131
domains 104
driver 11, 81
driver configuration file 86
driver installation 11
driver level 5

- driver package 155
- driver support 109
- driver utility 18
- driver version 4, 15
- drivers 6, 95, 139
- dust 122
- Dynix/ptx 107–108
- Dynix/ptx V4.4.10 108, 126
- Dynix/ptx V4.5.3 108, 117
- Dynix/ptx V4.5.x 126
- Dynix/ptx V4.6.1 108

E

- E3000 60
- EFC Manager 26, 102
- Emulex 2, 4–5, 11, 16, 35, 62, 81–82, 95, 109, 141, 155
- Emulex driver 17, 82
- Emulex LP6000 109
- Emulex LP8000 67, 95, 141
- Emulex LP8000DC 95, 141
- Emulex LP850 141
- Emulex LP9000 95, 141
- Emulex LP9002DC 95, 141
- Emulex LP9002L 95, 141
- Emulex LP9402DC 141
- Emulex LP952L 141
- Emulex tools 157
- enemy 116
- Enterprise Storage Server 60
- ES-1000 9, 99, 145
- eServer 33, 140
- ESS 119, 122, 125, 128, 138
- existing adapters 2
- extended distance 6, 130
- Extended LUN support 4
- EZ Fibre Configuration Utility 22

F

- F_Port 4, 103, 137, 151, 160
- fabric 96, 107, 110, 123, 129, 142
- fabric attachment 160
- fabric information 124
- fabric installation 117
- fabric levels 108
- fabric-attached 97, 131
- FASTT200 114, 146
- FASTT500 146

- FASTT700 146
- FC-AL 11, 36, 60, 95–96, 142
- fiber 7, 130
- fiber optic cable 6
- Fibre Channel adapter manufacturer 2
- Fibre Channel Arbitrated Loop 36
- filesystem 105
- firmware 5–6, 11, 15, 112, 129, 139
- firmware package 117
- FL_Port 4

G

- GBIC 6, 36, 96–97, 122, 130, 143
- Gigabit Interface Converters 6, 36, 96, 130, 143
- Gigabit Link Module 110
- GLM 110, 122

H

- HBA configuration files 65
- headers 147
- heterogeneous 113
- high speed communication 7, 96, 130, 143
- host adapter placement 111
- host bus adapter 1
- hot-pluggable adapters 101
- HP 9000 93, 95, 99
- Hp-ux 94, 100

I

- i386 139
- IEEE address 159
- independently configured 9
- infodev 80
- INRANGE 6–7, 38, 71, 96–97, 102, 142, 144, 150, 159
- INRANGE FC/9000 8, 98
- INRANGE FC/9000 40U 144
- INRANGE FC/9000 40U cabinet 98
- Intel 2
- Intel-based 140
- interactions 7, 97, 130, 143
- inter-switch links 6, 96, 143
- IN-VSN Enterprise Manager 26, 102, 159
- IO cards 108
- IP support 4
- ISL 6, 172

J

- JNI 2, 6, 11, 22, 35, 62, 81, 95, 141
- JNI 1160 21
- JNI driver 81
- JNI drivers 142
- JNI FC64-1063 63, 66
- JNI FCE2-6410 95, 141
- JNI FCE2-6412 95, 141
- JNI FCE-3210 141
- JNI FCE-6410 95
- JNI FCE-6460 64, 66, 77
- JNI FCI-1063 63, 66

K

- Kernel 147, 155
- kernel 108, 125, 141–142, 147
- kernel messages 149, 152, 157, 161
- kernel source files 147
- kernel-headers 147, 155
- kernel-source 147, 155
- KGPSA-BC 132
- KGPSA-CA 132

L

- L1000 101
- LAN 117
- large scale applications 5
- LEDs 17
- Light Pulse 17–18
- LiLO 155
- limitations 96
- link attenuation 7, 37, 96, 130, 143
- Linux 34, 139, 146
- Linux kernel 142
- loadable kernel 148
- loadable kernel modules 156
- login 136
- longwave 6, 96–98, 112, 130–131, 143
- loop attached devices 73
- loop switch 9
- loop-attached 97, 131
- LP6000 111
- LP7000E 110
- LP8000 4, 17, 62, 67, 155
- LP9000 4, 17
- LTO 3584 10
- LUN 124–125, 153, 161
- LUN masking 7, 97, 130, 143

- LUN's 119

- LUNs 122

M

- Magstar 3590 drives 10
- mapping storage 23
- McDATA 6, 8–9, 38, 72, 96, 98, 102, 142, 145, 150, 159
- McDATA FC-512 145
- McDATA FC-512 Cabinet 98
- McDATA switch 99, 145
- memory 108
- meshed 114
- Microsoft certified 4
- module dependencies 149
- monolithic kernel 147
- multimode 97–98, 110, 112, 128, 131
- multi-path 128
- multipath driver 108
- multipathing driver 60
- multiple protocols 5
- multi-switch 114

N

- N_Port 4, 102
- names 14
- Netfinity 11
- networking 108
- Node WWN 122
- NUMA-Q 107, 113
- NUMA-Q zone 113

O

- oil 122
- open source 140
- operating system level 5

P

- partitions 153
- patch 132
- patches 61, 77, 95
- path redundancy 125
- paths 108, 124
- PCI 11, 100, 111, 132–133
- PDC 100
- performance 96
- physical location 116

Point to Point 108
Point-To-Point 11, 129
Point-to-Point 5
point-to-point 142
Port World Wide Name 102
power 17
price 2
probe 123
processor building blocks 108
Processor Dependent Code 100
protocol 102
pSeries 33
PTTOPT_FABRIC 102

Q

QLA2200 12
QLA2200F 11, 64, 68, 147
QLA2300 12
QLA2300F 11
QLogic 2–4, 11–12, 15, 62, 81, 141, 147, 152
QLogic configuration 13
QLogic driver 82
QLogic Fast Utility 16
QLogic QLA2200F 68
QLogic SANblade 2200 141
QLogic SANblade 2300 141
QLview 13
QLview for Fibre Setup 13
QLview stand-alone configuration 13
QLview utility 13
quad 108, 110–111
quad0 118
quads 108

R

Red Hat 140
Red Hat 7.1 146
Red Hat Linux 147, 155
Redbooks Web site 180
 Contact us xviii
Request for Price Quotation 116
RPQ 116
RS/6000 33

S

SAN Data Gateway 114
SAN switch 15

SAN switches 2
SANblade 11
SANblade 2200 3
SANblade 2300 3
SC 94, 117, 129
scalability 96
SCSI 114, 161
SCSI bridge 114
SCSI LUNs 150
scsi-qla0-adapter-port 150
SDD 60, 108
SDG 114
seated 96
security 7, 97, 130, 143
Service Support Representative 101
SES 126
shortwave 6, 96–98, 110, 112, 128–131, 143
Silkworm 112
singlemode 97–98, 112, 131
single-path 34, 128
single-processor kernel 156
SMIT 42
Solaris 65
Solaris 2.6 60
Solaris 7 60
Solaris 8 60
SP 33
Sparc5 60
SSR 101
standard layer product installation model 112
statistics 13
storage adapter 4
storage devices 13
storage systems 9
StorageWorks 128
StorWatch 122
Subsystem Device Driver 60, 94, 128
Sun 450 76
Sun Microsystems 59
Sun Solaris 59
SuSE 140
SuSE 7.2 146
Switch 108
switch 113, 127
switch firmware 121
switch information 136
Switched Fabric 5, 11
Switched fabric 95
switched fabric 102

switchShow 79–80, 136
system requirement 6

T

TCP/IP 117
technical specifications 5–6
telnet 120–121
topologies 2, 108
TotalStorage 10
TotalStorage 2105 ESS 10
TotalStorage 2106 MSS 10
TotalStorage Enterprise Storage Server 146
TotalStorage FAStT200 10
TotalStorage FAStT500 10
TotalStorage FAStT700 10
Tru64 127
Tru64 Unix 128–129, 131–132, 136
TurboLinux 140

U

Ultra2 60
unbalanced fabric 125

V

VC-Encoded Address Mode 126
version 79
volume 151
volume group 105

W

Windows 1
Windows 2000 2, 11, 16
Windows boot 16
Windows NT 2, 11–12, 17
wizard 12
World Wide Name 102
WWN 15, 76, 102–103, 119, 122–123, 125, 135,
137, 150–151, 160

Z

zone 74, 113
zoned 108
zones 103, 151, 160
zoning 7, 37, 74, 97, 130, 143



Introducing Hosts to the SAN Fabric

(0.2" spine)
0.17" x 0.473"
90 x 249 pages



Introducing Hosts to the SAN Fabric



Ready a switched fabric capable host for SAN connectivity

Discover the simplicity of becoming SANitized

Verify your SANity in a complex world

In this IBM Redbook, we aim to be short, sharp, and to the point — and provide an answer. There is no marketing, no sales pitch, no extraneous information, other than what is needed to get the job done. That job is making hosts Storage Area Network (SAN) ready.

To this end, we show how to make hosts SAN fabric ready and connect them to the fabric with minimum fuss. We cover all SAN switched fabric capable hosts.

With the vast combination of SAN components, system levels, hardware and software, it is not possible to cover each and every environment that is likely to be encountered in the field. However, we do intend to cover those elements that are common, and that any reader with an intermediate grasp of SAN and host concepts, should be able to apply.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks