

IBM TotalStorage DS8000



Host Systems Attachment Guide

IBM TotalStorage DS8000



Host Systems Attachment Guide

Note:

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

Fourth Edition (May 2005)

This edition replaces SC26-7628-02 and all previous versions of SC26-7628.

© **Copyright International Business Machines Corporation 2004, 2005. All rights reserved.**

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

Contents

Figures	xiii
Tables	xv
About this guide	xvii
Safety and environmental notices	xix
Safety notices	xix
Environmental notices	xix
Product recycling	xix
Disposing of products	xix
Conventions used in this guide	xix
Related information	xx
DS8000 series library	xx
Other IBM publications	xxi
Ordering IBM publications	xxvi
Web sites	xxvi
How to send your comments	xxvii
 Summary of Changes for SC26-7628-03 IBM TotalStorage DS8000 Host Systems Attachment Guide	 xxix
 Chapter 1. Introduction to IBM TotalStorage DS8000	 1
DS8000 models	1
DS8100 (Model 921)	4
DS8300 (Models 922 and 9A2)	5
Model comparison	5
DS8000 physical footprint	6
DS8000 performance features	8
Interfaces of the DS8000	8
IBM TotalStorage DS Storage Manager	9
DS Open application programming interface	9
DS command-line interface	9
Path considerations for a storage image	10
FlashCopy, Remote Mirror and Copy attachment restrictions	10
IBM TotalStorage Multipath Subsystem Device Driver	11
DS8000 hardware specifics	11
Storage unit structure	11
IBM TotalStorage Management Console	12
Host systems that the storage unit supports	13
Attaching a DS8000 series to an open-systems host with fibre channel adapters	13
Fibre-channel (SCSI-FCP) attached open-systems hosts that the storage unit supports	14
ESCON-attached S/390 and zSeries hosts that the storage unit supports	15
FICON-attached S/390 and zSeries hosts that the storage unit supports	15
Processor memory	16
General information about attaching to open-systems host with fibre-channel adapters	16
Fibre-channel architecture	16
Fibre-channel cables and adapter types	19
Fibre-channel node-to-node distances	19
LUN affinity for fibre-channel attachment	20

Targets and LUNs for fibre-channel attachment	20
LUN access modes for fibre-channel attachment	20
Fibre-channel storage area networks	21
Data management features	22
RAID	22
Storage System LPARs (logical partitions)	22
Copy Services	23
FlashCopy	25
Disaster recovery using Copy Services	26
Comparison of licensed functions	27
Parallel Access Volumes	28
Chapter 2. Attaching to a Apple Macintosh server	29
Supported fibre-channel adapters for the Apple Macintosh server	29
Chapter 3. Attaching to a Fujitsu PRIMEPOWER host system	31
Supported fibre-channel adapters for PRIMEPOWER	31
Fibre-channel attachment requirements for PRIMEPOWER	31
Installing the Emulex adapter card for a PRIMEPOWER host system	31
Downloading the Emulex adapter driver for a PRIMEPOWER host system	32
Installing the Emulex adapter driver for a PRIMEPOWER host system	32
Configuring host device drivers for PRIMEPOWER	32
Parameter settings for the Emulex LP9002L adapter	34
Setting parameters for Emulex adapters	36
Chapter 4. Attaching to a Hewlett-Packard AlphaServer Tru64 UNIX host	37
Attaching to an HP AlphaServer Tru64 UNIX host system with fibre-channel adapters	37
Supported fibre-channel adapters for the HP AlphaServer Tru64 UNIX host system	37
Supported operating system levels for fibre-channel attachment to an HP Tru64 UNIX host	37
Fibre-channel Tru64 UNIX attachment requirements	37
Fibre-channel Tru64 UNIX attachment considerations.	38
Supporting the AlphaServer Console for Tru64 UNIX	38
Attaching the HP AlphaServer Tru64 UNIX host to a storage unit using fibre-channel adapters	39
Installing the KGPSA-xx adapter card in an Tru64 UNIX host system	40
Setting the mode for the KGPSA-xx host adapter	41
Setting up the storage unit to attach to an HP AlphaServer Tru64 UNIX host system with fibre-channel adapters.	42
Configuring the storage for fibre-channel Tru64 UNIX hosts	46
Removing persistent reserves for Tru64 UNIX 5.x	47
Limitations for Tru64 UNIX	50
Chapter 5. Attaching to a Hewlett-Packard AlphaServer OpenVMS host	51
Supported fibre-channel adapters for the HP AlphaServer OpenVMS host system	51
Fibre-channel OpenVMS attachment requirements.	51
Fibre-channel OpenVMS attachment considerations	51
Supported OpenVMS feature codes	52
Supported microcode levels for the HP OpenVMS host system	52
Supported switches for the HP OpenVMS host system	52
Attaching the HP AlphaServer OpenVMS host to a storage unit using fibre-channel adapters	52

Supported operating system levels for fibre-channel attachment to an HP OpenVMS host	52
Confirming the installation of the OpenVMS operating system.	53
Installing the KGPSA-xx adapter card in an OpenVMS host system	53
Setting the mode for the KGPSA-xx host adapter in an OpenVMS host system	54
Setting up the storage unit to attach to an HP AlphaServer OpenVMS host system with fibre-channel adapters.	55
Adding or modifying AlphaServer fibre-channel connections for the OpenVMS host	55
Defining OpenVMS fibre-channel adapters to the storage unit.	55
Configuring fibre-channel host adapter ports for OpenVMS.	56
OpenVMS fibre-channel considerations	56
OpenVMS UDID Support	56
OpenVMS LUN 0 - Command Control LUN	58
Confirming fibre-channel switch connectivity for OpenVMS.	58
Confirming fibre-channel storage connectivity for OpenVMS	59
OpenVMS World Wide Node Name hexadecimal representations	60
Verifying the fibre-channel attachment of the storage unit volumes for OpenVMS	60
Configuring the storage for fibre-channel OpenVMS hosts	61
OpenVMS fibre-channel restrictions	61
Troubleshooting fibre-channel attached volumes for the OpenVMS host system	62
Chapter 6. Attaching to a Hewlett-Packard Servers (HP-UX) host	65
Attaching with fibre-channel adapters.	65
Supported fibre-channel adapters for HP-UX hosts.	65
Fibre-channel attachment requirements for HP-UX hosts	65
Installing the fibre-channel adapter drivers for HP-UX 11.i, and HP-UX 11iv2	66
Setting the queue depth for the HP-UX operating system with fibre-channel adapters	66
Configuring the storage unit for clustering on the HP-UX 11iv2 operating system	66
Chapter 7. Attaching to an IBM iSeries host	69
Attaching with fibre-channel adapters to the IBM iSeries host system	69
Supported fibre-channel adapter cards for IBM iSeries hosts	69
Fibre-channel attachment requirements for IBM iSeries hosts.	69
Fibre-channel attachment considerations for IBM iSeries hosts	69
Host limitations for IBM iSeries hosts.	70
IBM iSeries hardware	71
IBM iSeries software	71
General information for configurations for IBM iSeries hosts	71
Recommended configurations for IBM iSeries hosts	72
Running the Linux operating system on an IBM i5 server	74
Supported fibre-channel adapters for IBM i5 servers running the Linux operating system	74
Running the Linux operating system in a guest partition on an IBM i5 servers	74
Planning to run Linux in a hosted or nonhosted guest partition	75
Creating a guest partition to run Linux	76
Managing Linux in a guest partition	77
Ordering a new server or upgrading an existing server to run a guest partition	78
Chapter 8. Attaching to an IBM NAS Gateway 500 host	79
Supported adapter cards for IBM NAS Gateway 500 hosts.	79
Finding the worldwide port name	79
Obtaining WWPNS using a Web browser	79
Obtaining WWPNS through the command-line interface	79

Multipathing support for NAS Gateway 500	81
Multipath I/O and SDD considerations for NAS Gateway 500	81
Host attachment multipathing scripts	81
Multipathing with the Subsystem Device Driver	83
Chapter 9. Attaching to an IBM RS/6000 or IBM eServer pSeries host	85
Installing the 2107 host attachment package on IBM pSeries AIX hosts	85
Preparing for installation of the host attachment package on IBM pSeries AIX hosts.	85
Installing the host attachment package on IBM pSeries AIX hosts	85
Upgrading the host attachment package on IBM pSeries AIX hosts.	86
Attaching with fibre-channel adapters.	86
Supported fibre-channel adapter cards for IBM pSeries hosts	86
Fibre-channel attachment requirements for IBM pSeries hosts	87
Fibre-channel attachment considerations for IBM pSeries hosts	87
Verifying the configuration of the storage unit for fibre-channel adapters on the AIX host system	87
Making SAN changes for IBM pSeries hosts	88
Support for fibre-channel boot	89
Prerequisites for setting up the IBM pSeries host as a fibre-channel boot device	89
Fibre-channel boot considerations for IBM pSeries hosts	89
Supported IBM RS/6000 or IBM pSeries hosts for fibre-channel boot	89
Supported levels of firmware for fibre-channel boot on IBM pSeries hosts	90
Supported levels of fibre-channel adapter microcode on IBM pSeries hosts	90
Installation mechanisms that PSSP supports for boot install from fibre-channel SAN DASD on IBM pSeries hosts	90
Support for disk configurations for RS/6000 for a fibre-channel boot install	90
Support for fibre-channel boot when a disk subsystem is attached on IBM pSeries hosts	90
Attaching to multiple RS/6000 or pSeries hosts without the HACMP host system	91
Considerations for attaching to multiple RS/6000 or pSeries hosts without the HACMP host system	91
Saving data on the storage unit when attaching multiple RS/6000 or pSeries host systems to the storage unit.	91
Restoring data on the storage unit when attaching multiple RS/6000 or pSeries host systems to the storage unit	92
Running the Linux operating system on an IBM pSeries host	92
Attachment considerations for running the Linux operating system on an IBM pSeries host	92
Hardware requirements for the Linux operating system on the pSeries host	93
Software requirement for the Linux operating system on the pSeries host	93
Preparing to install the Subsystem Device Driver for the Linux operating system on the pSeries host	94
Installing the Subsystem Device Driver on the pSeries host running the Linux operating system	94
Upgrading the Subsystem Device Driver for the Linux operating system on the pSeries host	95
Verifying the Subsystem Device Driver for the Linux operating system on the pSeries host	95
Configuring the Subsystem Device Driver	96
Migrating with AIX 5L Version 5.2	96
Storage unit migration issues when you upgrade to AIX 5L Version 5.2 maintenance release 5200-01	96
Storage unit migration issues when you remove the AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package support.	98

Chapter 10. Attaching to an IBM S/390 or IBM eServer zSeries host	99
ESCON connectivity on zSeries hosts	99
Attaching with ESCON adapters	100
Controller images and interconnections for ESCON adapters on zSeries hosts	100
Host adapters, cables, distances, and specifications for ESCON adapters on zSeries hosts	100
Cable lengths and types for ESCON adapters on zSeries hosts	101
Logical paths and path groups for ESCON adapters on zSeries hosts	101
Data transfer for ESCON adapters on zSeries hosts.	101
Directors and channel extenders for ESCON adapters on zSeries hosts	101
Identifying the port for TSO commands on zSeries hosts	101
Migrating from ESCON to native FICON	102
Migrating from ESCON to native FICON on zSeries hosts: FICON support	102
Migrating from ESCON to native FICON on zSeries hosts: Native ESCON configuration	102
Migrating from ESCON to native FICON on zSeries hosts: Mixed configuration	103
Migrating from ESCON to native FICON on zSeries hosts: FICON configuration	104
Migrating from a FICON bridge to a native FICON attachment	105
Migrating from a FICON bridge to a native FICON attachment on zSeries hosts: FICON bridge overview	105
Migrating from a FICON bridge to a native FICON attachment on zSeries hosts: FICON bridge configuration	105
Migrating from a FICON bridge to a native FICON attachment on zSeries hosts: Mixed configuration	106
Migrating from a FICON bridge to a native FICON attachment on zSeries hosts: Native FICON configuration	107
Attaching with FICON adapters	108
Configuring the storage unit for FICON attachment on zSeries hosts.	108
Attachment considerations for attaching with FICON adapters on zSeries hosts	108
Attaching to a FICON channel on a S/390 or zSeries host	110
Registered state-change notifications (RSCNs) on zSeries hosts	111
Linux for S/390 and zSeries.	112
Attaching a storage unit to an S/390 or zSeries host running Linux	112
Running Linux on an S/390 or zSeries host	112
Attaching FCP adapters on zSeries hosts.	113
 Chapter 11. Attaching to an Intel host running Linux	 119
Supported adapter cards for an Intel host running Linux	119
Attaching with fibre-channel adapters	119
Attachment requirements for an Intel host running Linux	119
Attachment considerations for an Intel host running Linux.	120
Installing the Emulex adapter card for an Intel host running Linux.	120
Downloading the Emulex adapter driver for an Intel host running Linux	121
Installing the Emulex adapter driver for an Intel host running Linux	121
Installing the QLogic adapter card for an Intel host running Linux	122
Downloading the current QLogic adapter driver for an Intel host running Linux	123
Installing the QLogic adapter driver for an Intel host running Linux	123
Defining the number of disk devices on Linux	123
Configuring the storage unit.	125
Configuring the storage unit for an Intel host running Linux	125
Partitioning storage unit disks for an Intel host running Linux	125

Assigning the system ID to the partition for an Intel host running Linux	126
Creating and using file systems on the storage unit for an Intel host running Linux	127
SCSI disk considerations for an Intel host running Linux	128
Manually adding and removing SCSI disks	128
LUN identification for the Linux host system	128
SCSI disk problem identification and resolution	130
Support for fibre-channel boot	130
Creating a modules disk for SUSE LINUX Enterprise Server 9.0	130
Installing Linux over the SAN without an IBM Subsystem Device Driver	130
Updating to a more recent module without an IBM Subsystem Device Driver	132
Installing Linux over the SAN with an IBM Subsystem Device Driver	132
Chapter 12. Attaching to an Intel host running VMware ESX server	139
Supported adapter cards for an Intel host running VMware ESX Server	139
Attaching with fibre-channel adapters	139
Attachment requirements for an Intel host running VMware ESX server	139
Attachment considerations for an Intel host running VMware ESX Server	140
Installing the Emulex adapter card for an Intel host running VMware ESX Server	140
Installing the QLogic adapter card for an Intel host running VMware ESX Server	141
Defining the number of disks devices on VMware ESX Server	142
SCSI disk considerations for an Intel host running VMware ESX server	142
LUN identification for the VMware ESX console	142
Disk device discovery on VMware ESX	145
Persistent binding	145
Configuring the storage unit	146
Configuring the storage unit for an Intel host running VMware ESX Server	146
Partitioning storage unit disks for an Intel host running VMware ESX Server	146
Creating and using VMFS on the storage unit for an Intel host running VMware ESX Server	148
Copy Services considerations	148
Chapter 13. Attaching to a Novell NetWare host	151
Attaching with fibre-channel adapters	151
Attaching a Novell NetWare host with fibre-channel adapters	151
Fibre-channel attachment considerations for a Novell NetWare host	151
Installing the Emulex adapter card for a Novell NetWare host	151
Downloading the current Emulex adapter driver for a Novell NetWare host	152
Installing the Emulex adapter driver for a Novell NetWare host	152
Downloading the current QLogic adapter driver for a Novell NetWare host	152
Installing the QLogic QLA23xx adapter card for a Novell NetWare host	153
Installing the adapter drivers for a Novell NetWare host	154
Chapter 14. Attaching to a iSCSI Gateway host	157
Attachment considerations	157
Attachment overview of the iSCSI Gateway host	157
Attachment requirements for the iSCSI Gateway host	157
Ethernet adapter attachment considerations for the iSCSI Gateway host	158
Configuring for storage for the iSCSI Gateway host	158
iSCSI Gateway operation through the IP Service Module	159
Chapter 15. Attaching to an IBM SAN File System	161
Attaching to an IBM SAN File System metadata server with Fibre-channel adapters	161

Configuring a storage unit for attachment to the SAN File System metadata node	161
Chapter 16. Attaching to an IBM SAN Volume Controller host	163
Attaching to an IBM SAN Volume Controller host with fibre-channel adapters	163
Configuring the storage unit for attachment to the SAN Volume Controller host system	163
Chapter 17. Attaching an SGI host system	165
Attaching an SGI host system with fibre-channel adapters	165
Fibre-channel attachment considerations for the SGI host system	165
Fibre-channel attachment requirements for the SGI host system	165
Checking the version of the IRIX operating system for the SGI host system	166
Installing a fibre-channel adapter card for the SGI host system	166
Verifying the installation of a fibre-channel adapter card for SGI	166
Configuring the fibre-channel adapter drivers for SGI	167
Installing an optical cable for SGI in a switched-fabric topology	167
Installing an optical cable for SGI in an arbitrated loop topology	167
Confirming switch connectivity for SGI	167
Displaying zoning information for the switch	168
Confirming storage connectivity	169
Confirming storage connectivity for SGI in a switched-fabric topology	169
Confirming storage connectivity for SGI in a fibre-channel arbitrated loop topology	170
Configuring the storage unit for host failover	170
Configuring the storage unit for host failover	170
Confirming the availability of failover	171
Making a connection through a switched-fabric topology	171
Making a connection through an arbitrated-loop topology	172
Switching I/O operations between the primary and secondary paths	172
Configuring storage	172
Configuration considerations	172
Configuring storage in a switched fabric topology	173
Configuring storage in an arbitrated loop topology	175
Chapter 18. Attaching to a Sun host	179
Attaching with fibre-channel adapters	179
Supported fibre-channel adapters for Sun	179
Fibre-channel attachment requirements for Sun	179
Installing the Emulex adapter card for a Sun host system	180
Downloading the Emulex adapter driver for a Sun host system	181
Installing the Emulex adapter driver for a Sun host system	181
Installing the AMCC PCI adapter card for Sun	181
Downloading the current AMCC PCI adapter driver for Sun	182
Installing the AMCC PCI adapter driver for Sun	182
Installing the AMCC SBUS adapter card for Sun	182
Downloading the current AMCC SBUS adapter driver for Sun	182
Installing the AMCC SBUS adapter driver for Sun	182
Installing the QLogic adapter card in the Sun host	183
Downloading the QLogic adapter driver for Sun	183
Installing the QLogic adapter driver package for Sun	183
Configuring host device drivers	184
Configuring host device drivers for Sun	184
Parameter settings for the Emulex adapters for the Sun host system	185
Parameter settings for the AMCC adapters for the Sun host system	187
Parameter settings for the QLogic QLA23xxF adapters	190

Parameter settings for the QLogic QLA23xx adapters for San Surf configuration (4.06+ driver)	191
Setting the Sun host system parameters	192
Setting parameters for AMCC adapters	193
Setting parameters for Emulex or QLogic adapters	193
Installing the Subsystem Device Driver	194
Attaching a storage unit to a Sun host using Storage Traffic Manager System	194
Configuring the Sun STMS host settings	195
Attaching a storage unit to a Sun host using Sun Cluster	195
Chapter 19. Attaching to a Windows 2000 host	197
Attaching with fibre-channel adapters	197
Supported fibre-channel adapters for Windows 2000	197
Fibre-channel attachment requirements for the Windows 2000 host system	197
Fibre-channel attachment considerations for Windows 2000	198
Installing and configuring the Emulex adapter card	198
Installing and configuring the Netfinity adapter card for Windows 2000	201
Updating the Windows 2000 device driver	203
Installing and configuring the QLogic adapter cards in Windows 2000	203
Verifying that Windows 2000 is configured for storage	206
Configuring for availability and recoverability	206
Configuration considerations	206
Setting the TimeOutValue registry for Windows 2000	206
Installing remote fibre-channel boot support for a Windows 2000 host system	207
Configure zoning and obtain storage	207
Flash QLogic host adapter	208
Configure QLogic host adapters	208
Windows 2000 Installation	208
Windows 2000 Post Installation	209
Chapter 20. Attaching to a Windows Server 2003 host.	211
Attaching with fibre-channel adapters	211
Supported fibre-channel adapters for Windows Server 2003	211
Fibre-channel attachment requirements for the Windows Server 2003 host system.	211
Fibre-channel attachment considerations for Windows Server 2003	212
Installing and configuring the Emulex adapter card	212
Installing and configuring the Netfinity adapter card	215
Updating the Windows Server 2003 device driver.	217
Installing and configuring the QLogic adapter cards	217
Verifying that Windows Server 2003 is configured for storage	220
Configuring for availability and recoverability	220
Configuration considerations	220
Setting the TimeOutValue registry for Windows Server 2003.	220
Installing remote fibre-channel boot support for a Windows Server 2003 host system	221
Configure zoning and obtain storage	221
Flash QLogic host adapter	222
Configure QLogic host adapters	222
Windows 2003 Installation	222
Windows 2003 Post Installation	223
Appendix. Locating the worldwide port name (WWPN)	225
Fibre-channel port name identification	225
Locating the WWPN for a Fujitsu PRIMEPOWER host.	225
Locating the WWPN for a Hewlett-Packard AlphaServer host	226

Locating the WWPN for a Hewlett-Packard host	226
Locating the WWPN for an IBM eServer iSeries host	227
Locating the WWPN for an IBM eServer pSeries or an RS/6000 host	227
Locating the WWPN for a Linux host	228
Locating the WWPN for a Novell NetWare host	229
Locating the WWPN for an SGI host	229
Locating the WWPN for a Sun host	230
Locating the WWPN for a Windows 2000 host	230
Locating the WWPN for a Windows Server 2003 host	231
Accessibility	233
Notices	235
Terms and conditions for downloading and printing publications	236
Trademarks.	237
Electronic emission notices	238
Federal Communications Commission (FCC) statement	238
Industry Canada compliance statement	238
European community compliance statement.	238
Japanese Voluntary Control Council for Interference (VCCI) class A statement	239
Korean Ministry of Information and Communication (MIC) statement.	239
Taiwan class A compliance statement	240
Glossary	241
.	241
Index	269

Figures

1. A base model (with front cover off) and its main components	2
2. An expansion model (with front cover off) and its main components	3
3. Maximum configuration for Model 921	4
4. Maximum configuration for Models 922 and 9A2.	5
5. DS8000 physical footprint. Dimensions are in centimeters (inches)..	7
6. Point-to-point topology	17
7. Switched-fabric topology	18
8. Arbitrated loop topology	19
9. Example of sd.conf file entries for fibre-channel	33
10. Example of a start lpfc auto-generated configuration	34
11. Confirming the storage unit licensed machine code on an HP AlphaServer through the telnet command	39
12. Example of the sizer -v command	40
13. Example of the set mode diag command and the wwidmgr -show adapter command	41
14. Example results of the wwidmgr command.	41
15. Example of the switchshow command	43
16. Example of storage unit volumes on the AlphaServer console	43
17. Example of a hex string for a storage unit volume on an AlphaServer Tru64 UNIX console	44
18. Example of a hex string that identifies the decimal volume number for a storage unit volume on an AlphaServer console or Tru64 UNIX	44
19. Example of hex representation of last 5 characters of a storage unit volume serial number on an AlphaServer console	44
20. Example of the hwmgr command to verify attachment	45
21. Example of a Korn shell script to display a summary of storage unit volumes	46
22. Example of the Korn shell script output.	46
23. Example of the essvol script	48
24. Example of the scu command	48
25. Example of the scu command	49
26. Example of the scu clear command	49
27. Example of the scu command to show persistent reserves	50
28. Example of show system command to show system command on the OpenVMS operating system	53
29. Example of the product show history command to check the versions of patches already installed	53
30. Example of the set mode diag command and the wwidmgr -show adapter command	54
31. Example results of the wwidmgr command.	55
32. Example of the switchshow command	59
33. Example of storage unit volumes on the AlphaServer console	59
34. Example of a World Wide Node Name for the storage unit volume on an AlphaServer console	60
35. Example of a volume number for the storage unit volume on an AlphaServer console	60
36. Example of what is displayed when you use OpenVMS storage configuration utilities.	61
37. Example of the display for the auxiliary storage hardware resource detail for a 2766 or 2787 adapter card	70
38. Example of the logical hardware resources associated with an IOP	72
39. Example of the display for the auxiliary storage hardware resource detail for the storage unit	72
40. Example output from the lscfg -vpl "fcs*" grep Network command.	80
41. Example output saved to a text file	80
42. Example of installed file set	82
43. Example environment where all MPIO devices removed	82
44. Example of installed host attachment scripts.	83
45. Example of extracting the SDD archive.	84
46. Example of a list of devices displayed when you use the lsdev -Cc disk grep 2107 command for fibre-channel	88
47. Example of a list of other devices displayed when you use the lsdisk command for fibre-channel.	88
48. ESCON connectivity	99

49. Port identification for S/390 and zSeries TSO commands	102
50. Example of an ESCON configuration	103
51. Example of an ESCON configuration with added FICON channels	104
52. Example of a native FICON configuration with FICON channels that have been moved nondisruptively	105
53. Example of how to configure a FICON bridge from an S/390 or zSeries host system to a storage unit	106
54. Example of how to add a FICON director and a FICON host adapter	107
55. Example of the configuration after the FICON bridge is removed	108
56. Example of prerequisite information for FCP Linux on zSeries	114
57. Example of prerequisite information for FCP Linux on zSeries	115
58. Example of prerequisite information for FCP Linux on zSeries	116
59. Example of prerequisite information for FCP Linux on zSeries	117
60. Example of a script to add more than one device	117
61. Example of how to add SCSI devices through the add_map command	118
62. Saving the module parameters in the /etc/zfcp.conf file	118
63. Example of Logical Volume Manager Multipathing	118
64. Example of range of devices for a Linux host	125
65. Example of different options for the fdisk utility	126
66. Example of a primary partition on the disk /dev/sdb.	126
67. Example of assigning a Linux system ID to the partition	127
68. Example of creating a file with the mke2fs command	127
69. Example of creating a file with the mkfs command	128
70. Example output for a Linux host that only configures LUN 0	129
71. Example output for a Linux host with configured LUNs	129
72. Example of a complete linuxrc file for Red Hat	137
73. Example of a complete linuxrc file for SUSE	138
74. Example of QLogic Output:	143
75. Example of Emulex Output:	144
76. Example listing of a Vmhba directory	144
77. Example of Vmhba entries	145
78. Example of the different options for the fdisk utility:	147
79. Example of primary partition on the disk /dev/vsd71	147
80. Example of PCI bus slots	167
81. Example results for the switchshow command	168
82. Example results for the cfgShow command.	169
83. Example of commands to turn failover on	171
84. Example of an edited /etc/failover.conf file	172
85. Example of an edited /etc/failover.conf file for an arbitrated loop connection.	172
86. Example commands for the IRIX switched fabric storage configuration utility	174
87. Example commands for the IRIX switched fabric storage configuration utility, part 2	175
88. Example commands for the IRIX arbitrated loop storage configuration utility	176
89. Example commands for the IRIX arbitrated loop storage configuration utility, part 2	177
90. Example of sd.conf file entries for fibre-channel	184
91. Example of a start lpfc auto-generated configuration	185
92. Example binding inserts for qlaxxxx.conf.	192
93. Example of what is displayed when start the Windows 2000 host	202
94. Example of what the system displays when you start the Windows Server 2003 host	216
95. Example of the output from the Hewlett-Packard AlphaServer wwidmgr -show command	226
96. Example of the output from the Hewlett-Packard #fgrep wwn /var/adm/messages command	226
97. Example of the output from the Hewlett-Packard fcmsutil /dev/td1 grep world command.	227
98. Example of what is displayed in the /var/log/messages file	229
99. Example of the scsiha — bus_number device command	229

Tables

1. DS8000 model comparison chart	6
2. Comparison of licensed functions	27
3. Recommended configuration file parameters for the Emulex LP9002L adapter	35
4. Maximum number of adapters you can use for an AlphaServer	38
5. Maximum number of adapters you can use for an AlphaServer	51
6. Host system limitations for the IBM iSeries host system	70
7. Capacity and models of disk volumes for IBM iSeries	73
8. Recommended settings for the QLogic adapter card for an Intel host running Linux	122
9. Recommended settings for the QLogic adapter card for an Intel host running VMware ESX Server	141
10. Recommended settings for the QLogic QLA23xx adapter card for a Novell NetWare host	153
11. Solaris 8 minimum revision level patches for fibre-channel	180
12. Recommended configuration file parameters for the Emulex LP9002DC, LP9002L, LP9002S, LP9402DC and LP9802 adapters	186
13. Recommended configuration file parameters for a AMCC FCX-6562, AMCC FCX2-6562, AMCC FCE-6460, or a AMCC FCE-1473 adapter	187
14. Recommended configuration file parameters for the QLogic QLA2310F, QLA2340, and QLA2342 adapters with driver level 4.03	190
15. Parameter settings for the QLogic QLA23xx host adapters for San Surf Configuration (4.06+)	191
16. Recommended configuration file parameters for the Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, and LP10000DC adapters	200
17. Recommended settings for the QLogic QLA23xx adapter card for Windows 2000	204
18. Recommended configuration file parameters for the Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, and LP10000DC adapters	214
19. Recommended settings for the QLogic QLA23xx adapter card for a Windows Server 2003 host	218

About this guide

This guide provides information about the following host attachment issues:

- Attaching the storage unit to an open-systems host with fibre-channel adapters
- Connecting IBM Enterprise Systems Connection (ESCON) cables to your IBM S/390 and IBM eServer zSeries host systems
- Connecting IBM fibre-channel connection (FICON) cables to your S/390 and zSeries host systems

You can attach the following host systems to a storage unit:

- Apple Macintosh
- Fujitsu PRIMEPOWER
- Hewlett-Packard
- IBM eServer iSeries
- IBM NAS Gateway 500
- IBM RS/6000 and IBM eServer pSeries
- IBM S/390 and IBM eServer zSeries
- IBM SAN File System
- IBM SAN Volume Controller
- Linux
- Microsoft Windows 2000
- Microsoft Windows Server 2003
- Novell NetWare
- Silicon Graphics
- Sun

For a list of open systems hosts, operating systems, adapters and switches that IBM supports, see the *Interoperability Matrix* at

<http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Finding attachment information in the Host Systems Attachment Guide

The following table provides information about how to find attachment information quickly in this guide.

Host	Fibre-channel
Apple Macintosh Servers	"Supported fibre-channel adapters for the Apple Macintosh server" on page 29
Fujitsu PRIMEPOWER	"Supported fibre-channel adapters for PRIMEPOWER" on page 31
Hewlett-Packard AlphaServer Tru64 UNIX®	"Supported fibre-channel adapters for the HP AlphaServer Tru64 UNIX host system" on page 37
Hewlett-Packard AlphaServer OpenVMS	"Supported fibre-channel adapters for the HP AlphaServer OpenVMS host system" on page 51
Hewlett-Packard Servers (HP-UX)	"Attaching with fibre-channel adapters" on page 65

Host	Fibre-channel
iSCSI Gateway	"Attachment overview of the iSCSI Gateway host" on page 157
IBM® eServer™ iSeries™	"Attaching with fibre-channel adapters to the IBM iSeries host system" on page 69
IBM NAS Gateway 500	"Supported adapter cards for IBM NAS Gateway 500 hosts" on page 79
IBM eServer pSeries™ or IBM RS/6000®	"Attaching with fibre-channel adapters" on page 86
IBM eServer zSeries® or S/390®	"Attaching FCP adapters on zSeries hosts running Linux" on page 113 (for zSeries only)
Linux	"Attaching with fibre-channel adapters" on page 119
Microsoft® Windows® 2000	"Attaching with fibre-channel adapters" on page 197
Microsoft Windows Server 2003	"Attaching with fibre-channel adapters" on page 211
Novell NetWare	"Attaching with fibre-channel adapters" on page 151
IBM SAN Volume Controller	"Attaching to an IBM SAN Volume Controller host with fibre-channel adapters" on page 163
Silicon Graphics	"Attaching an SGI host system with fibre-channel adapters" on page 165
Sun	"Supported fibre-channel adapters for Sun" on page 179

Safety and environmental notices

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

Safety notices

Use this process to find information about safety notices.

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

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

DANGER

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

1000

CAUTION:

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

1001

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

Environmental notices

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

Product recycling

This unit contains recyclable materials.

Recycle these materials at your local recycling sites. Recycle the materials according to local regulations. In some areas, IBM provides a product take-back program that ensures proper handling of the product. Contact your IBM representative for more information.

Disposing of products

This topic contains information about how to dispose of products.

This unit might contain batteries. Remove and discard these batteries, or recycle them, according to local regulations.

Conventions used in this guide

The following typefaces are used to show emphasis:

boldface

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

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

monospace

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

Related information

The tables in this section list and describe the following publications:

- The publications that make up the IBM® TotalStorage™ DS8000 series library
- Other IBM publications that relate to the DS8000 series
- Non-IBM publications that relate to the DS8000 series

See “Ordering IBM publications” on page xxvi for information about how to order publications in the IBM TotalStorage DS8000 series publication library. See “How to send your comments” on page xxvii for information about how to send comments about the publications.

DS8000 series library

These customer publications make up the DS8000 series library.

Unless otherwise noted, these publications are available in Adobe portable document format (PDF) on a compact disc (CD) that comes with the storage unit. If you need additional copies of this CD, the order number is SK2T-8803. These publications are also available as PDF files by clicking on the **Documentation link** on the following Web site:

<http://www-1.ibm.com/servers/storage/support/disk/ds8100/index.html>

See “Ordering IBM publications” on page xxvi for information about ordering these and other IBM publications.

Title	Description	Order Number
<i>IBM TotalStorage® DS: Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the command-line interface (CLI) for managing your DS8000 configuration and Copy Services relationships. The CLI application provides a set of commands that you can use to write customized scripts for a host system. The scripts initiate predefined tasks in a Copy Services server application. You can use the CLI commands to indirectly control Remote Mirror and Copy and FlashCopy® configuration tasks within a Copy Services server group.	SC26-7625 (See Note.)
<i>IBM TotalStorage DS8000: Host Systems Attachment Guide</i>	This guide provides guidelines for attaching the DS8000 to your host system and for migrating to fibre-channel attachment from a small computer system interface.	SC26-7628 (See Note.)
<i>IBM TotalStorage DS8000: Introduction and Planning Guide</i>	This guide introduces the DS8000 product and lists the features you can order. It also provides guidelines for planning the installation and configuration of the storage unit.	GC35-0495

Title	Description	Order Number
<i>IBM TotalStorage Multipath Subsystem Device Driver User's Guide</i>	This publication describes how to use the IBM Subsystem Device Driver (SDD) on open-systems hosts to enhance performance and availability on the DS8000. SDD creates redundant paths for shared logical unit numbers. SDD permits applications to run without interruption when path errors occur. It balances the workload across paths, and it transparently integrates with applications.	SC30-4096
<i>IBM TotalStorage DS8000: User's Guide</i>	This guide provides instructions for setting up and operating the DS8000 and for analyzing problems.	SC26-7623 (See Note.)
<i>IBM TotalStorage DS Application Programming Interface Reference</i>	This publication provides reference information for the IBM TotalStorage DS application programming interface (API) and provides instructions for installing the Common Information Model Agent, which implements the API.	GC35-0493
<i>IBM TotalStorage DS8000 Messages Reference</i>	This publication provides explanations of error, information, and warning messages that are issued from the DS8000 user interfaces.	GC26-7659
Note: No hardcopy book is produced for this publication. However, a PDF file is available from the following Web site: http://www-1.ibm.com/servers/storage/support/disk/ds8100/index.html		

Other IBM publications

Other IBM publications contain additional information that is related to the DS product library.

The following list is divided into categories to help you find publications that are related to specific topics. Some of the publications are listed under more than one category. See "Ordering IBM publications" on page xxvi for information about ordering these and other IBM publications.

Title	Description	Order Number
Data-copy services		
<i>z/OS DFSMS Advanced Copy Services</i>	This publication helps you understand and use IBM Advanced Copy Services functions. It describes three dynamic copy functions and several point-in-time copy functions. These functions provide backup and recovery of data if a disaster occurs to your data center. The dynamic copy functions are peer-to-peer remote copy, extended remote copy, and coupled extended remote copy. Collectively, these functions are known as remote copy. FlashCopy, SnapShot, and concurrent copy are the point-in-time copy functions.	SC35-0428
<i>IBM Enterprise Storage Server</i>	This publication, from the IBM International Technical Support Organization, introduces the Enterprise Storage Server and provides an understanding of its benefits. It also describes in detail the architecture, hardware, and functions, including the advanced copy functions, of the Enterprise Storage Server.	SG24-5465
<i>Implementing Copy Services On S/390</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an Enterprise Storage Server that is attached to an S/390 or zSeries host system. Copy Services functions include peer-to-peer remote copy, extended remote copy, FlashCopy®, and concurrent copy. This publication describes the functions, prerequisites, and corequisites and describes how to implement each function into your environment.	SG24-5680

Title	Description	Order Number
<i>IBM TotalStorage ESS Implementing Copy Services in an Open Environment</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on UNIX, Windows NT®, Windows 2000, Sun Solaris, HP-UX, Tru64, OpenVMS, and iSeries host systems. The Copy Services functions that are described include peer-to-peer remote copy and FlashCopy. This publication describes the functions and shows you how to implement them into your environment. It also shows you how to implement these functions in a high-availability cluster multiprocessing environment.	SG24-5757
Fibre channel		
<i>Fibre Channel Connection (FICON) I/O Interface: Physical Layer</i>	This publication provides information about the fibre-channel I/O interface. This book is also available as a PDF file from the following Web site: http://www.ibm.com/servers/resourcelink/	SA24-7172
<i>Fibre Transport Services (FTS): Physical and Configuration Planning Guide</i>	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide</i>	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.	SC26-7350
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide</i>	This guide describes the IBM SAN Fibre Channel Switch and the IBM TotalStorage ESS Specialist. It provides information about the commands and how to manage the switch with Telnet and the Simple Network Management Protocol.	SC26-7349
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide</i>	This publication describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S16. It is intended for trained service representatives and service providers.	SC26-7352
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide</i>	This guide introduces the IBM SAN Fibre Channel Switch 2109 Model S16 and tells you how to manage and monitor the switch using zoning and how to manage the switch remotely.	SC26-7351
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the Enterprise Storage Server. It provides you with a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment.	SG24-6113
Open-systems hosts		
<i>ESS Solutions for Open Systems Storage: Compaq AlphaServer, HP, and Sun</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure the Enterprise Storage Server when you attach Compaq AlphaServer (running Tru64 UNIX), HP, and Sun hosts. This book does not cover Compaq AlphaServer that is running the OpenVMS operating system. This book also focuses on the settings that are required to give optimal performance and on the settings for device driver levels. This book is for the experienced UNIX professional who has a broad understanding of storage concepts.	SG24-6119

Title	Description	Order Number
<i>IBM TotalStorage ESS Implementing Copy Services in an Open Environment</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on UNIX or Windows 2000 host systems. The Copy Services functions that are described include peer-to-peer remote copy and FlashCopy. This publication describes the functions and shows you how to implement them into your environment. It also shows you how to implement these functions in a high-availability cluster multiprocessing environment.	SG24-5757
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the Enterprise Storage Server. It gives you a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment.	SG24-6113
S/390 and zSeries hosts		
<i>Device Support Facilities: User's Guide and Reference</i>	This publication describes the IBM Device Support Facilities (ICKDSF) product that are used with IBM direct access storage device (DASD) subsystems. ICKDSF is a program that you can use to perform functions that are needed for the installation, the use, and the maintenance of IBM DASD. You can also use it to perform service functions, error detection, and media maintenance.	GC35-0033
<i>z/OS Advanced Copy Services</i>	This publication helps you understand and use IBM Advanced Copy Services functions. It describes three dynamic copy functions and several point-in-time copy functions. These functions provide backup and recovery of data if a disaster occurs to your data center. The dynamic copy functions are peer-to-peer remote copy, extended remote copy, and coupled extended remote copy. Collectively, these functions are known as remote copy. FlashCopy, SnapShot, and concurrent copy are the point-in-time copy functions.	SC35-0428
<i>DFSMS/MVS V1: Remote Copy Guide and Reference</i>	This publication provides guidelines for using remote copy functions with S/390 and zSeries hosts.	SC35-0169
<i>Fibre Transport Services (FTS): Physical and Configuration Planning Guide</i>	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234
<i>Implementing ESS Copy Services on S/390</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an Enterprise Storage Server that is attached to an S/390 or zSeries host system. Copy Services functions include peer-to-peer remote copy, extended remote copy, FlashCopy, and concurrent copy. This publication describes the functions, prerequisites, and corequisites and describes how to implement each function into your environment.	SG24-5680
<i>ES/9000, ES/3090: IOCP User Guide Volume A04</i>	This publication describes the Input/Output Configuration Program that supports the Enterprise Systems Connection (ESCON) architecture. It describes how to define, install, and configure the channels or channel paths, control units, and I/O devices on the ES/9000 processors and the IBM ES/3090 Processor Complex.	GC38-0097
<i>IOCP User's Guide, IBM e(logo)server zSeries 800 and 900</i>	This publication describes the Input/Output Configuration Program that supports the zSeries 800 and 900 servers. This publication is available in PDF format by accessing ResourceLink at the following Web site: www.ibm.com/servers/resourceLink/	SB10-7029

Title	Description	Order Number
<i>IOCP User's Guide, IBM e(logo)server zSeries</i>	This publication describes the Input/Output Configuration Program that supports the zSeries server. This publication is available in PDF format by accessing ResourceLink at the following Web site: www.ibm.com/servers/resourceLink/	SB10-7037
<i>S/390: Input/Output Configuration Program User's Guide and ESCON Channel-to-Channel Reference</i>	This publication describes the Input/Output Configuration Program that supports ESCON architecture and the ESCON multiple image facility.	GC38-0401
<i>IBM z/OS Hardware Configuration Definition User's Guide</i>	This guide provides conceptual and procedural information to help you use the z/OS Hardware Configuration Definition (HCD) application. It also explains: <ul style="list-style-type: none"> • How to migrate existing IOCP/MVSCP definitions • How to use HCD to dynamically activate a new configuration • How to resolve problems in conjunction with MVS/ESA HCD 	SC33-7988
<i>OS/390: Hardware Configuration Definition User's Guide</i>	This guide provides detailed information about the input/output definition file and about how to configure parallel access volumes. This guide discusses how to use Hardware Configuration Definition for both OS/390® and z/OS V1R1.	SC28-1848
<i>OS/390 V2R10.0: MVS System Messages Volume 1 (ABA - ASA)</i>	This publication lists OS/390 MVS™ system messages ABA to ASA.	GC28-1784
<i>Using IBM 3390 Direct Access Storage in a VM Environment</i>	This publication provides device-specific information for the various models of the 3390 and describes methods you can use to manage storage efficiently using the VM operating system. It provides guidance on managing system performance, availability, and space through effective use of the direct access storage subsystem.	GG26-4575
<i>Using IBM 3390 Direct Access Storage in a VSE Environment</i>	This publication helps you use the 3390 in a VSE environment. It includes planning information for adding new 3390 units and instructions for installing devices, migrating data, and performing ongoing storage management activities.	GC26-4576
<i>Using IBM 3390 Direct Access Storage in an MVS Environment</i>	This publication helps you use the 3390 in an MVS environment. It includes device-specific information for the various models of the 3390 and illustrates techniques for more efficient storage management. It also offers guidance on managing system performance, availability, and space utilization through effective use of the direct access storage subsystem.	GC26-4574
<i>z/Architecture Principles of Operation</i>	This publication provides a detailed definition of the z/Architecture™. It is written as a reference for use primarily by assembler language programmers and describes each function at the level of detail needed to prepare an assembler language program that relies on a particular function. However, anyone concerned with the functional details of z/Architecture will find this publication useful.	SA22-7832
SAN		

Title	Description	Order Number
<i>IBM OS/390 Hardware Configuration Definition User's Guide</i>	This guide explains how to use the Hardware Configuration Data application to perform the following tasks: <ul style="list-style-type: none"> • Define new hardware configurations • View and modify existing hardware configurations • Activate configurations • Query supported hardware • Maintain input/output definition files (IODFs) • Compare two IODFs or compare an IODF with an actual configuration • Print reports of configurations • Create graphical reports of a configuration • Migrate existing configuration data 	SC28-1848
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide</i>	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.	SC26-7350
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide</i>	This guide describes the IBM SAN Fibre Channel Switch and the IBM TotalStorage ESS Specialist. It provides information about the commands and how to manage the switch with Telnet and the Simple Network Management Protocol (SNMP).	SC26-7349
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide</i>	This publication describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S16. It is intended for trained service representatives and service providers.	SC26-7352
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide</i>	This guide introduces the IBM SAN Fibre Channel Switch 2109 Model S16 and tells you how to manage and monitor the switch using zoning and how to manage the switch remotely.	SC26-7351
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the Enterprise Storage Server. It provides you with a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment.	SG24-6113
Seascope family		
<i>IBM Enterprise Storage Server</i>	This publication, from the IBM International Technical Support Organization, introduces the Enterprise Storage Server and provides an understanding of its benefits. It also describes in detail the architecture, hardware, and functions, including the advanced copy functions, of the Enterprise Storage Server.	SG24-5465
<i>IBM Enterprise Storage Server Performance Monitoring and Tuning Guide</i>	This guide, from the IBM International Technical Support Organization, provides guidance on the best way to configure, monitor, and manage your Enterprise Storage Server to ensure optimum performance.	SG24-5656
<i>IBM Versatile Storage Server: Introduction and Planning Guide</i>	This publication introduces the IBM Versatile Storage Server™ and lists the features you can order. It also provides planning information for both 2105 Models B09 and 100.	GC26-7223
<i>Implementing the IBM Enterprise Storage Server in Your Environment</i>	This publication, from the IBM International Technical Support Organization, can help you install, tailor, and configure the Enterprise Storage Server in your environment.	SG24-5420
Storage management		

Title	Description	Order Number
<i>Device Support Facilities: User's Guide and Reference</i>	This publication describes the IBM Device Support Facilities (ICKDSF) product used with IBM direct access storage device (DASD) subsystems. ICKDSF is a program that you can use to perform functions that are needed for the installation, the use, and the maintenance of IBM DASD. You can also use it to perform service functions, error detection, and media maintenance.	GC35-0033
<i>IBM TotalStorage Solutions Handbook</i>	This handbook, from the IBM International Technical Support Organization, helps you understand what makes up enterprise storage management. The concepts include the key technologies that you must know and the IBM subsystems, software, and solutions that are available today. It also provides guidelines for implementing various enterprise storage administration tasks so that you can establish your own enterprise storage management environment.	SG24-5250

Ordering IBM publications

This section tells you how to order copies of IBM publications and how to set up a profile to receive notifications about new or changed publications.

IBM publications center

The publications center is a worldwide central repository for IBM product publications and marketing material.

The IBM publications center offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download free of charge. You can also order publications. The publications center displays prices in your local currency. You can access the IBM publications center through the following Web site:

<http://www.ibm.com/shop/publications/order>

Publications notification system

The IBM publications center Web site offers you a notification system for IBM publications.

If you register, you can create your own profile of publications that interest you. The publications notification system sends you a daily e-mail that contains information about new or revised publications that are based on your profile.

If you want to subscribe, you can access the publications notification system from the IBM publications center at the following Web site:

<http://www.ibm.com/shop/publications/order>

Web sites

The following Web sites provide information about the IBM TotalStorage DS8000 series and other IBM storage products.

Type of Storage Information	Web Site
Concurrent Copy for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/

Type of Storage Information	Web Site
Copy Services command-line interface (CLI)	http://www-1.ibm.com/servers/storage/support/software/cscli.html
DS8000 series publications	http://www-1.ibm.com/servers/storage/support/disk/ds8100/index.html Click Documentation .
FlashCopy for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
Host system models, operating systems, and adapters that the storage unit supports	http://www.ibm.com/servers/storage/disk/ds8000/interop.html Click Interoperability matrix .
IBM Disk Storage Feature Activation (DSFA)	http://www.ibm.com/storage/dsfa
IBM storage products	http://www.storage.ibm.com/
IBM TotalStorage DS8000 series	http://www-1.ibm.com/servers/storage/disk/ds8000
IBM version of the Java (JRE) that is often required for IBM products	http://www-106.ibm.com/developerworks/java/jdk/
Multiple Device Manager (MDM)	http://www.ibm.com/servers/storage/support/ Click Storage Virtualization .
Remote Mirror and Copy (formerly PPRC) for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
SAN fibre channel switches	http://www.ibm.com/storage/fcswitch/
Storage Area Network Gateway and Router	http://www-1.ibm.com/servers/storage/support/san/index.html
Subsystem Device Driver (SDD)	http://www-1.ibm.com/servers/storage/support/software/sdd.html
z/OS Global Mirror (formerly XRC) for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this information or any other DS8000 series documentation, you can submit them in the following ways:

- e-mail

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

starpubs@us.ibm.com

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

- Mail

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

International Business Machines Corporation
RCF Processing Department

Department 61C
9032 South Rita Road
TUCSON AZ 85775-4401

Summary of Changes for SC26-7628-03 IBM TotalStorage DS8000 Host Systems Attachment Guide

This document contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. This summary of changes describes new functions that have been added to this release.

Changed Information

This edition includes the following changed information:

- Support is added for SUSE SLES 9 and Red Hat Enterprise Linux 4.0 for select hosts.
- The Hewlett-Packard OpenVMS chapter DS CLI command example was updated.
- The URL to locate the most current information on host attachment firmware and device driver information, including driver downloads, changed to:

<http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do>

Chapter 1. Introduction to IBM TotalStorage DS8000

IBM TotalStorage DS8000 is a high-performance, high-capacity series of disk storage that is designed to support continuous operations. DS8000 series models (machine type 2107) use the IBM POWER5 server technology that is integrated with the IBM Virtualization Engine technology. DS8000 series models consist of a storage unit and one or two management consoles, two being the recommended configuration. The graphical user interface (GUI) or the command-line interface (CLI) allows you to logically partition storage and use the built-in Copy Services functions. For high-availability, hardware components are redundant.

DS8000 models

The DS8000 series offers various choices of base and expansion models, so you can configure storage units that meet your performance and configuration needs.

DS8100

The DS8100 (Model 921) features a dual two-way processor complex and support for one expansion frame.

DS8300

The DS8300 (Models 922 and 9A2) features a dual four-way processor complex and support for one or two expansion frames. The Model 9A2 supports two IBM TotalStorage System LPARs (Logical Partitions) in one storage unit.

The DS8000 expansion frames (Models 92E and 9AE) expand the capabilities of the base models. You can attach the Model 92E to either the Model 921 or the Model 922 to expand their capabilities. You can attach the Model 9AE to expand the Model 9A2.

Figure 1 on page 2 provides a high-level view at the components for the DS8100 and DS8300 base models (Models 921, 922, and 9A2).





Figure 1. A base model (with front cover off) and its main components

The following notes provide additional information about the labeled components on Figure 1:

1. The rack power area of the base models provides redundant power supplies (two primary power supplies), power control cards, and backup battery assemblies to help protect data in the event of a loss of external power. Model 921 contains two batteries. Model 922 and Model 9A2 contain three batteries to support the 4-way processors.
2. All base models allow up to eight disk enclosures, which contain the disk drives. In a maximum configuration, each base model can hold up to 128 disk drives.
3. All base models contain one management console, a keyboard, and display. The management console is the focal point for configuration, copy services management, and maintenance activities.
4. All base models contain two processor enclosures. The processor enclosures contain the computer electronic complex (CEC). The Model 921 processor enclosures have 2-way processors. Processor enclosures on the Model 922 and Model 9A2 have 4-way processors.
5. All base models contain I/O enclosures and adapters. The I/O enclosures hold the adapters and provide connectivity between the adapters and the processors. Both device adapters and host adapters are installed in the I/O enclosure.

Figure 2 on page 3 provides a high-level view at the components for the expansion models (Models 92E and 9AE).



1vxc86

Figure 2. An expansion model (with front cover off) and its main components

The following notes provide additional information about the labeled components on Figure 2:

1. The rack power area of each expansion model provides redundant power supplies (two primary power supplies) and power control cards. If the expansion unit contains I/O enclosures or the extended power line disturbance (PLD) feature, three backup battery assemblies are also installed to help protect data in the event of a loss of external power.
2. All expansion models allow up to 16 disk enclosures, which contain the disk drives. In a maximum configuration, each expansion model can hold 256 disk drives.
3. Expansion models can contain I/O enclosures and adaptors if they are the first expansion units that are attached to either a Model 922 or a Model 9A2. The second expansion model in a 922 or 9A2 configuration cannot have I/O enclosures and adaptors, nor can any expansion unit that is attached to a Model 921. If the expansion unit contains I/O enclosures, the enclosures provide connectivity between the adaptors and the processors. The adaptors contained in the I/O enclosures can be either device or host adaptors.

DS8100 (Model 921)

The IBM TotalStorage DS8100, which is Model 921, offers many features.

These features include:

- Dual two-way processor complex
- Up to 128 disk drives, for a maximum capacity of 38.4 TB
- Up to 128 GB of processor memory (cache)
- Up to 16 fibre-channel/FICON or ESCON® host adapters

The DS8100 model can support one expansion frame. With one expansion frame, you can expand the capacity of the Model 921 as follows:

- Up to 384 disk drives, for a maximum capacity of 115.2 TB

Note: IBM service representatives can upgrade a Model 921 in the field when you order a model conversion to a Model 922 or Model 9A2.

Figure 3 shows the maximum configuration of a Model 921, which is the 921 base model plus one 92E expansion model.

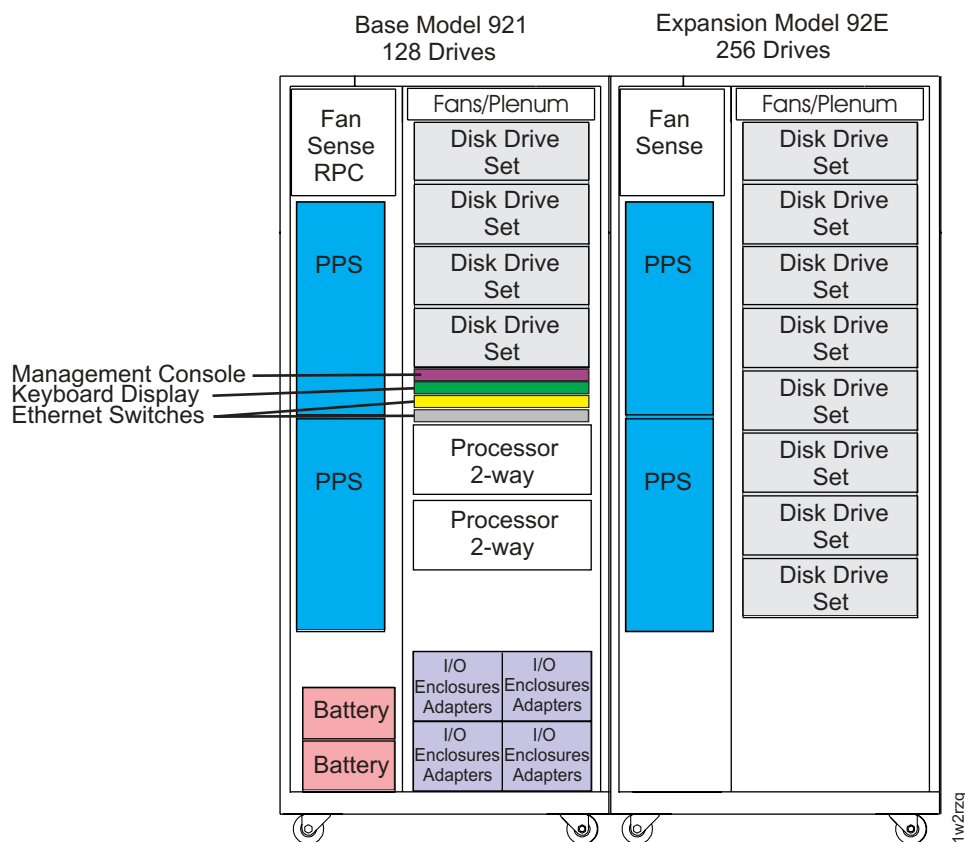


Figure 3. Maximum configuration for Model 921

DS8300 (Models 922 and 9A2)

IBM TotalStorage DS8300 models (Model 922 and Model 9A2) offer higher performance and capacity than the DS8100. The Model 9A2 also enables you to create two storage system LPARs (or images) within the same storage unit.

Both DS8300 models offer the following features:

- Dual four-way processor complex
- Up to 128 disk drives, for a maximum capacity of 38.4 TB
- Up to 256 GB of processor memory (cache)
- Up to 16 fibre-channel/FICON or ESCON host adapters

The DS8300 models can support either one or two expansion frames. With expansion frames, you can expand the Model 922 and 9A2 as follows:

- With one expansion frame, you can support the following expanded capacity and number of adapters:
 - Up to 384 disk drives, for a maximum capacity of 115.2 TB
 - Up to 32 fibre-channel/FICON or ESCON host adapters
- With two expansion frames, you can support the following expanded capacity:
 - Up to 640 disk drives, for a maximum capacity of 192 TB

Figure 4 shows the maximum configuration of a Model 922 and 9A2. Both of these models can attach up to two expansion models. Model 922 uses the 92E expansion model. Model 9A2 uses the 9AE expansion model.

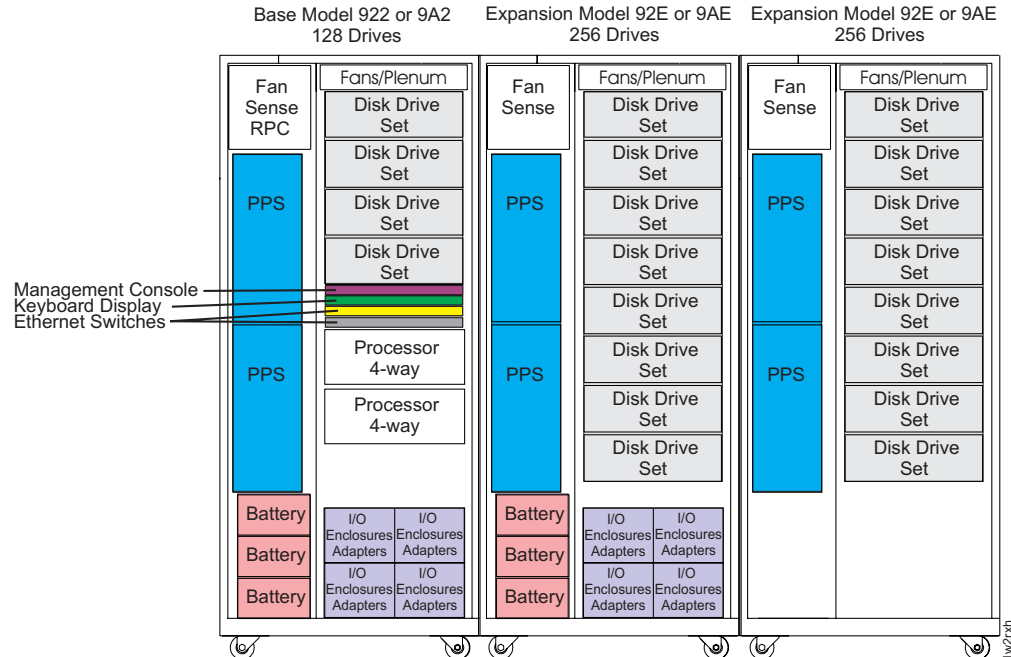


Figure 4. Maximum configuration for Models 922 and 9A2

Model comparison

DS8000 models vary on the processor types that they use, their disk capacity, and other supported features.

Table 1 compares the various supported features that are available on different DS8000 model configurations.

Table 1. DS8000 model comparison chart

Base model	Images (logical partitions)	Expansion models	Processor type	Disk drives	Processor memory	Host adapters
Model 921	1 image	None	2-way	Up to 128 disks (maximum capacity of 38.4 TB)	Up to 128 GB	Up to 16 host adapters
		With 1 expansion model (92E)	2-way	Up to 384 disks (maximum capacity of 115.2 TB)	Up to 128 GB	Up to 16 host adapters
Model 922	1 image	None	4-way	Up to 128 disks (maximum capacity of 38.4 TB)	Up to 256 GB	Up to 16 host adapters
		With 1 expansion model (92E)	4-way	Up to 384 disks (maximum capacity of 115.2 TB)	Up to 256 GB	Up to 32 host adapters
		With 2 expansion models (92E)	4-way	Up to 640 disks (maximum capacity of 192 TB)	Up to 256 GB	Up to 32 host adapters
Model 9A2	2 images	None	4-way	Up to 128 disks (maximum capacity of 38.4 TB)	Up to 256 GB	Up to 16 host adapters
		With 1 expansion model (9AE)	4-way	Up to 384 disks (maximum capacity of 115.2 TB)	Up to 256 GB	Up to 32 host adapters
		With 2 expansion models (9AE)	4-way	Up to 640 disks (maximum capacity of 192 TB)	Up to 256 GB	Up to 32 host adapters

Note: IBM service representatives can upgrade a Model 921 in the field when you order a model conversion to a Model 922 or a Model 9A2.

DS8000 physical footprint

The physical footprint dimensions, caster locations, and cable openings for a DS8000 unit help you plan your installation site.

[illegible]

The following dimensions are labeled on Figure 5:

- Chapter 1. Introduction to IBM TotalStorage DS8000
- 7**

DS8000 performance features

The DS8000 series is designed to provide you with high-performance, connectivity, and reliability allowing your workload to be easily consolidated into a single storage subsystem.

The following list provides an overview of some of the features that are associated with the DS8000 series:

POWER5 processor technology

The DS8000 features IBM POWER5 server technology. Depending on workload, the maximum host I/O operations per second of the DS8100 Model 921 is up to three times the maximum operations per second of the ESS Model 800. The maximum host I/O operations per second of the DS8300 Model 922 or 9A2 is up to six times the maximum of the ESS model 800.

Industry standard fibre channel disk drives

The DS8000 series offers a selection of fibre channel disk drives, including 300 GB drives, allowing a DS8100 to scale up to 115.2 TB of capacity and a DS8300 to scale up to 192 TB of capacity.

Four-port fibre channel/FICON adapters

These new adapters not only enhance connectivity, but increase configuration flexibility because the individual ports can be configured to support fibre channel or FICON.

HOST attachments

The DS8100 offers up to 16 host adapters (64 FCP/FICON ports) and the DS8300 offers up to 32 host adapters (128 FCP/FICON ports), further increasing your ability to share your disk storage.

IBM Standby Capacity on Demand

The IBM Standby Capacity on Demand (Standby CoD) offering allows the installation of inactive disk drives that can be easily activated as business needs require. With this offering, up to four Standby CoD disk drive sets (64 disk drives) can be factory or field installed into your system. To activate, you logically configure the disk drives for use—a non-disruptive activity that does not require intervention from IBM. Upon activation of any portion of the Standby CoD disk drive set, you must place an order with IBM to initiate billing for the activated set. At that time, you can also order replacement Standby CoD disk drive sets.

Online Information Center

The online Information Center is an information database that provides you the opportunity to quickly familiarize yourself with the major aspects of the DS8000 and to easily recognize the topics for which you might require more information. It provides information regarding user assistance for tasks, concepts, reference, user scenarios, tutorials, and other types of user information. Because the information is all in one place rather than across multiple publications, you can access the information that you need more efficiently and effectively.

Interfaces of the DS8000

The DS8000 provides management interfaces that contain management tools and utilities to help you increase productivity.

IBM TotalStorage DS Storage Manager

The IBM TotalStorage DS Storage Manager is a program interface that is used to perform logical configurations and Copy Services management functions.

The DS Storage Manager program is installed as a GUI (graphical mode) or as an unattended (silent mode) installation for the supported operating systems. It can be accessed from any location that has network access using a Web browser. It offers you the following choices:

Simulated configuration

You install this component on your PC or the Master Console which provides the ability to create or modify logical configurations when your storage unit is disconnected from the network. After creating the configuration, you can save it and then apply it to a network attached storage unit at a later time.

Real-time configuration

This component is preinstalled on your HMC. It provides you the ability to create logical configurations and use Copy Services features when your storage unit is attached to the network. This component provides you with real-time (online) configuration support.

DS Open application programming interface

The DS Open application programming interface (API) is a nonproprietary storage management client application that supports routine LUN management activities, such as LUN creation, mapping and masking, and the creation or deletion of RAID5 and RAID10 volume spaces. The DS Open API also enables Copy Services functions such as FlashCopy and Remote Mirror and Copy (formally known as peer-to-peer remote copy).

The IBM TotalStorage DS Open API helps integrate DS configuration management support into storage resource management (SRM) applications, which allow customers to benefit from existing SRM applications and infrastructures. The DS Open API also enables the automation of configuration management through customer-written applications. Either way, the DS Open API presents another option for managing storage units by complementing the use of the IBM TotalStorage DS Storage Manager web-based interface and the DS command-line interface.

You must implement the DS Open API through the IBM TotalStorage Common Information Model (CIM) agent, a middleware application that provides a CIM-compliant interface. The DS Open API uses the CIM technology to manage proprietary devices as open system devices through storage management applications. The DS Open API allows these storage management applications to communicate with a storage unit.

DS command-line interface

The IBM TotalStorage DS Command-Line Interface (CLI) enables open systems hosts to invoke and manage FlashCopy and Metro and Global Mirror functions through batch processes and scripts.

The command-line interface provides a full-function command set that allows you to check your storage unit configuration and perform specific application functions when necessary.

The following list highlights a few of the specific types of functions that you can perform with the DS command-line interface:

- Check and verify your storage unit configuration
- Check the current Copy Services configuration that is used by the storage unit
- Create new logical storage and Copy Services configuration settings
- Modify or delete logical storage and Copy Services configuration settings

Note: You cannot install the DS CLI on a Windows 64-bit operating system.

Path considerations for a storage image

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

When you want to attach multiple paths from a single host system to I/O ports on a host adapter of a storage image, the following considerations apply:

- Choose the attached I/O ports on different host adapters. For example, you want to spread the connections so that they are on independent I/O cards.
- Balance the connections between even and odd I/O enclosures (left and right enclosures with a rack). The even and odd enclosures share a power management boundary for certain recovery actions.
- Attach host adapters to I/O enclosures on separate RIO-G loop halves. The two I/O enclosures of each I/O enclosure pair (I1/I2 or I3/I4) are located on different RIO loop halves. You can place two host attachments, one in each of the two I/O enclosures of any I/O enclosure pair.

For example, on a Model 922, the following sets of four I/O enclosures are on separate RIO loop halves on both loops and they are the best selection for four host attachments:

- R1-I1, R1-I2, R1-I3, R1-I4
- R2-I1, R2-I2, R2-I3, R2-I4
- R1-I1, R1-I2, R2-I1, R2-I2
- R1-I3, R1-I4, R2-I3, R2-I4

FlashCopy, Remote Mirror and Copy attachment restrictions

Some attachment restrictions apply to FlashCopy, Remote Mirror and Copy processing.

The source and target volumes should be on different host systems when you copy a source volume to a target volume with FlashCopy or Remote Mirror and Copy. The source and target volumes should be on different host systems when you require concurrent read/write access of both volumes. A copy operation with the source and target volume on the same host system creates a target volume with the same identification as the source volume. The host system sees two identical volumes.

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

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

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

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

For more information on supported configurations, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>

IBM TotalStorage Multipath Subsystem Device Driver

The IBM TotalStorage Multipath Subsystem Device Driver provides redundant connections between a disk storage server and a host server in a storage unit.

You get redundancy with the IBM TotalStorage Multipath Subsystem Device Driver (SDD). The SDD resides in the host server with the native disk-device driver for the IBM storage unit. It uses redundant connections between disk storage server and host server in a storage unit to provide data availability and performance.

The SDD provides the following functions:

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

For more information about SDD, see <http://www-1.ibm.com/servers/storage/support/software/sdd.html>.

DS8000 hardware specifics

The DS8000 models offer a high degree of availability and performance through the use of redundant components that can be replaced while the system is operating. You can use the DS8000 models with a mix of different operating systems and clustered and nonclustered variants of the same operating systems.

Contributing to the high degree of availability and reliability are the structure of the DS8000 storage unit, the host systems it supports, and its processor memory and processor speeds.

Storage unit structure

The design of the storage unit, which contains the base model and the expansion models, contributes to the high degree of availability that is associated with the DS8000. The primary components that support high availability within the storage unit are the storage server, the processor complex, and the rack power control card.

The storage unit also has redundant fibre switches, with fibre fabric connections to both processor complexes. The entire power system at all levels is fully redundant.

There are redundant private LAN networks between each storage unit and the management consoles (even if there is only one management console).

Storage unit

A storage unit consists of a storage server and two or more storage (disk) enclosures that are packaged in one or more racks with associated power supplies, batteries, and cooling.

Storage server

A storage server consists of two processor complexes, two or more I/O enclosures, and a pair of rack power control cards.

Processor complex

A processor complex controls and manages the storage unit to perform the function of the storage server. The two processor complexes form a redundant pair such that if either processor complex fails, the remaining processor complex performs all storage server functions.

Rack power control card

A redundant pair of rack power control (RPC) cards coordinate the power management within the storage unit. The RPC cards are attached to the service processors in each processor complex, the primary power supplies in each rack, and indirectly to the fan/sense cards and storage enclosures in each rack.

Multi-path subsystem device driver

The IBM TotalStorage Multi-path subsystem device driver (SDD) provides load balancing and enhanced data availability capability in configurations with more than one I/O path between the host server and the DS8000 series system. Load balancing can help reduce or eliminate I/O bottlenecks that occur when many I/O operations are directed to common devices that are using the same I/O path. The SDD also helps eliminate a potential single point of failure by automatically rerouting I/O operations when a path failure occurs, thereby supporting enhanced data availability capability.

IBM TotalStorage Management Console

The IBM TotalStorage Management Console is the focal point for configuration, copy services management, and maintenance activities.

The Management Console is a dedicated workstation that is physically located (installed) inside your DS8100 and DS8300 and can automatically monitor the state of your system, notifying you and IBM when service is required. The Management Console can also be connected to your network to enable centralized management of your system using the IBM TotalStorage DS Command-Line Interface or the storage management software that uses the IBM TotalStorage DS Open API.

To provide continuous availability of customer access to the management console functions, having a second management console is recommended. This second management console can be provided in two ways:

- **External** (outside the 2107) - This console is generally installed in the customer-provided rack. It uses the same hardware as the internal management console.
- **Internal** (a second management console that resides within a second 2107) - This console can be cabled together with the primary internal management console of the first 2107.'

Host systems that the storage unit supports

The DS8000 storage unit provides a variety of host attachments so that you can consolidate storage capacity and workloads for open-systems hosts, S/390 hosts, and eServer zSeries hosts.

The DS8100 Model 921 supports a maximum of 16 host adapters and 4 device adapter pairs, and the DS8300 Models 922 and 9A2 supports a maximum of 32 host adapters and 8 device adapter pairs.

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

- Fibre-channel adapters, for support of fibre-channel protocol (FCP) and fibre connection (FICON[®]) protocol
- Enterprise Systems Connection Architecture[®] (ESCON) adapters

For fibre-channel attachments, you can establish zones. The zones must contain a single port that is attached to a system adapter with the desired number of ports that are attached to the storage unit. By establishing zones, you reduce the possibility of interactions between system adapters in switched configurations. You can establish the zones by using either of two zoning methods:

- Port number
- Worldwide port name (WWPN)

You can configure switch ports that are attached to the storage unit in more than one zone. This enables multiple system adapters to share access to the storage unit fibre-channel ports. Shared access to a storage unit fibre-channel port might come from host platforms that support a combination of bus adapter types and the operating systems.

Attaching a DS8000 series to an open-systems host with fibre channel adapters

You can attach a DS8000 series to an open-systems host with fibre-channel adapters.

Fibre channel is a 1 Gbps or 2 Gbps, full-duplex, serial communications technology to interconnect I/O devices and host systems that are separated by tens of kilometers.

The IBM TotalStorage DS8000 series supports SAN connections at 1 Gbps to 4 Gbps with 2 Gbps host bus adapters. The DS8000 series negotiates automatically and determines whether it is best to run at 1 Gbps link or 2 Gbps link. The IBM TotalStorage DS8000 series detects and operates at the higher link speed.

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

Fibre-channel (SCSI-FCP) attached open-systems hosts that the storage unit supports

You can attach the DS8000 storage unit to fibre-channel (SCSI-FCP) attached open-systems hosts.

Each storage unit fibre-channel adapter has four ports. Each port has a unique worldwide port name (WWPN). You can configure the port to operate with the SCSI-FCP upper-layer protocol. Shortwave adapters and longwave adapters are available on the storage unit.

Fibre-channel adapters for SCSI-FCP support provide the following configurations:

- A maximum of 64 host ports for DS8100 Model 921 and a maximum of 128 host ports for DS8300 Models 922 and 9A2
- A maximum of 509 host logins per fibre-channel port
- A maximum of 2000 N-port logins per storage image
- Access to all 63.7KB LUNs per target (one target per host adapter), depending on host type
- Either arbitrated loop, switched fabric, or point-to-point topologies

The storage unit supports the following host systems for shortwave fibre-channel attachment and longwave fibre-channel attachment:

- IBM eServer iSeries servers that run an IBM OS/400® operating system
- IBM eServer i5 servers that run an IBM i5/OS™ operating system
- IBM RS/6000, pSeries, eServer p5, RS/6000 SP™, and pSeries SP server that run an IBM AIX operating system
- IBM zSeries servers that run the Linux™ operating system
- IBM zSeries servers that run the z/VM® operating system
- Apple Macintosh servers
- Fujitsu PRIMEPOWER servers that run the Solaris operating system
- Hewlett-Packard servers that run an HP Tru64 UNIX operating system
- Hewlett-Packard servers that run an HP OpenVMS operating system
- Hewlett-Packard servers that run an HP-UX operating system
- Intel-based servers that run a Microsoft Windows 2000 operating system
- Intel-based servers that run a Microsoft Windows Server 2003 operating system for both 32-bit and 64-bit configurations
- Intel-based servers that run a Novell NetWare operating system
- IBM servers that run the IBM SAN Volume Controller operating system
- IBM servers that run the IBM SAN File System operating system
- iSCSI Gateway servers that run a Microsoft Windows 2000, Microsoft Windows 2003, or UNIX operating system
- Linux servers that run a Red Hat Linux, Red Flag Linux, Asianux, and SUSE Linux operating system
- SGI servers that run an IRIX operating system
- Sun servers that run a Solaris operating system

See the Interoperability Matrix at

<http://www.ibm.com/servers/storage/disk/ds8000/interop.html> for details about types, models, adapters, and the operating systems that the storage unit supports.

ESCON-attached S/390 and zSeries hosts that the storage unit supports

You can attach the DS8000 storage unit to the ESCON-attached S/390 and zSeries hosts.

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

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

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

Note: ESCON host channels limit the number of devices per channel to 1024. To fully access 4096 devices on a storage unit, it is necessary to connect a minimum of four ESCON host channels to the storage unit. You can access the devices through a switch to a single storage unit ESCON port. This method exposes four control-unit images (1024 devices) to each host channel.

The FICON bridge card in ESCON director 9032 Model 5 enables a FICON bridge channel to connect to ESCON host adapters in the storage unit. The FICON bridge architecture supports up to 16 384 devices per channel. This enables you to attach other control units or other storage units to the same host channel up to the limit that the host supports.

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

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

For details about models and the operating system versions and releases that the storage unit supports for these host systems, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

FICON-attached S/390 and zSeries hosts that the storage unit supports

You can attach the DS8000 storage unit to FICON-attached S/390 and zSeries hosts.

Each storage unit fibre-channel adapter has four ports. Each port has a unique world wide port name (WWPN). You can configure the port to operate with the FICON upper-layer protocol. When configured for FICON, the fibre-channel port supports connections to a maximum of 128 FICON hosts. On FICON, the

fibre-channel adapter can operate with fabric or point-to-point topologies. With fibre-channel adapters that are configured for FICON, the storage unit provides the following configurations:

- Either fabric or point-to-point topologies
- A maximum of 64 host ports for DS8100 Model 921 and a maximum of 128 host ports for DS8300 Models 922 and 9A2
- A maximum of 509 channel connections per fibre-channel port
- A maximum of 2048 logical paths on each fibre-channel port
- Access to all 32 control-unit images (8000 CKD devices) over each FICON port

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

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

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

For details about models, versions of operating systems, and releases that the storage unit supports for these host systems, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Processor memory

The DS8100 Model 921 offers up to 128 GB of processor memory and the DS8300 Models 922 and 9A2 offer up to 256 GB of processor memory. The Non-Volatile Storage (NVS) scales to the selected processor memory size, which can also help optimize performance.

General information about attaching to open-systems host with fibre-channel adapters

You can attach a storage unit to host systems with fibre channel-adapters.

Fibre-channel architecture

Fibre-channel architecture provides communications protocols on the storage unit.

The storage unit provides a fibre-channel connection in the storage unit. For more information about hosts and operating systems that the storage unit supports on the fibre-channel adapters, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Fibre-channel architecture provides a variety of communication protocols on the storage unit. The units that are interconnected are referred to as *nodes*. Each node has one or more ports.

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

Storage unit architecture supports these basic interconnection topologies (network structure):

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

Point-to-point topology

The point-to-point topology, also known as direct connect, enables you to interconnect ports directly. Figure 6 shows an illustration of a point-to-point topology.



Legend

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

Figure 6. Point-to-point topology

The storage unit supports direct point-to-point topology at a maximum distance of 500 m (1500 ft) at 1 Gb and 300 m (984 ft) at 2 Gb with the shortwave adapter. The storage unit supports direct point-to-point topology at a maximum distance of 10 km (6.2 mi) with the longwave adapter.


Switched-fabric topology

The switched-fabric topology provides the underlying structure that enables you to interconnect multiple nodes. You can use a fabric that provides the necessary switching functions to support communication between multiple nodes.

You can extend the distance that the storage unit supports up to 300 km (186.3 miles) with a storage area network (SAN) or other fabric components.

The storage unit supports increased connectivity with the use of fibre-channel (SCSI-FCP and FICON) directors. Specific details on status, availability, and configuration options that are supported by the storage unit are available on <http://www-1.ibm.com/servers/storage/disk/ds8000>.

The storage unit supports the switched-fabric topology with point-to-point protocol. You should configure the storage unit fibre-channel adapter to operate in point-to-point mode when you connect it to a fabric topology. See Figure 7 on page 18.



Legend

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

Figure 7. Switched-fabric topology


Arbitrated loop topology

Fibre Channel-Arbitrated Loop (FC-AL) is a ring topology that enables you to interconnect a set of nodes. The maximum number of ports that you can have on an FC-AL is 127. See Figure 8 on page 19.

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

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

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



Legend

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

Figure 8. Arbitrated loop topology

Note: IBM supports only the topologies for point-to-point and arbitrated loop. Unconfigure the port to change the topology.

Fibre-channel cables and adapter types

This section provides information about fibre-channel cables and adapter types.

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

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

See *IBM TotalStorage DS8000 Introduction and Planning Guide* for detailed information about fibre-channel cables and adapter types. This document also includes information about cable features and optional cables.

Fibre-channel node-to-node distances

The storage unit supports fibre-channel adapter for extended node-to-node distances.

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

For fibre-channel, the maximum distance between the following items is 11 km (6 mi):

- Fabric switches
- Link extenders
- Host fibre-channel port

- Storage unit fibre-channel port

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

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

LUN affinity for fibre-channel attachment

You can use a WWPN to associate a LUN for fibre-channel attachment.

For fibre-channel attachment, LUNs have an affinity to the host's fibre-channel adapter through the worldwide port name (WWPN) for the host adapter. In a switched fabric configuration, a single fibre-channel host adapter could have physical access to multiple fibre-channel ports on the storage unit. In this case, you can configure the storage unit to allow the host to use either some or all of the physically accessible fibre-channel ports on the storage unit.

Targets and LUNs for fibre-channel attachment

You can attach a fibre-channel host adapter to a LUN.

For fibre-channel attachment, each fibre-channel host adapter can architecturally attach up to 2 LUNs. The storage unit supports 64K LUNs and a maximum LUN size of 2TB divided into a maximum of 16 logical subsystems each with up to 256 LUNs. You can configure no more than 256 of the LUNs in the storage unit to be accessible by that host.

LUN access modes for fibre-channel attachment

Several modes allow LUN access for fibre-channel attachment.

The following sections describe the LUN access modes for fibre-channel.

Fibre-channel access modes

The fibre-channel architecture allows any fibre-channel initiator to access any fibre-channel device, without access restrictions. However, in some environments this kind of flexibility can represent a security exposure. Therefore, the IBM TotalStorage DS8000 allows you to restrict this type of access when IBM sets the access mode for your storage unit during initial configuration. There are two types of LUN access modes:

1. Access-any mode

The access-any mode allows all fibre-channel attached host systems that do not have an access profile to access all non-iSeries open system logical volumes that you have defined in the storage unit.

Note: If you connect the storage unit to more than one host system with multiple platforms and use the access-any mode without setting up an access profile for the hosts, the data in the LUN used by one open-systems host might be inadvertently corrupted by a second open-systems host. Certain host operating systems insist on overwriting specific LUN tracks during the LUN discovery phase of the operating system start process.

2. Access-restricted mode

The access-restricted mode prevents all fibre-channel-attached host systems that do not have an access profile to access volumes that you defined in the storage unit. This is the default mode.

Your IBM service support representative (SSR) can change the logical unit number (LUN) access mode. Changing the access mode is a disruptive process. Shut down and restart both clusters of the storage unit.

Access profiles

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

The setup of an access profile is transparent to you when you use the IBM TotalStorage DS Storage Manager to configure the hosts and volumes in the storage unit. Configuration actions that affect the access profile are as follows:

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

The anonymous host

When you run the storage unit in access-any mode, the IBM TotalStorage DS Storage Manager displays a dynamically created pseudo-host called anonymous. This is not a real host system connected to the storage server. It is intended to represent all fibre-channel-attached host systems that are connected to the storage unit that do not have an access profile defined. This is a reminder that logical volumes defined in the storage unit can be accessed by hosts which have not been identified to the storage unit.

Fibre-channel storage area networks

Fibre-channel storage area networks connect servers and storage devices.

A Fibre-channel storage area network (SAN) is a specialized, high-speed network that attaches servers and storage devices. With a SAN, you can perform an any-to-any connection across the network using interconnect elements such as routers, gateways and switches. With a SAN, you can eliminate the connection between a server, storage, and the concept that the server effectively owns and manages the storage devices.

The SAN also eliminates any restriction to the amount of data that a server can access. This restriction is limited by the number of storage devices that can be attached to the individual server. Instead, a SAN introduces the flexibility of

networking to enable one server or many heterogeneous servers to share a common storage utility. This might comprise many storage devices, including disk, tape, and optical storage. You can locate the storage utility far from the servers that use it.

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

Fibre-channel SANs, however, provide the capability to interconnect open systems and storage in the same network as S/390 and zSeries host systems and storage. You can map the protocols for attaching open systems and S/390 and zSeries host systems to the FC-4 layer of the fibre-channel architecture.

Data management features

The DS8000 storage unit is designed with the following management features that allow you to securely process and access your data according to your business needs even if it is 24 hours a day and 7 days a week.

RAID

Redundant array of independent disks (RAID) is a method of configuring multiple disk drives in a storage subsystem for high availability and high performance.

The collection of two or more disk drives presents the image of a single disk drive to the system. In the event of a single device failure, data can be read or regenerated from the other disk drives in the array. With RAID implementation, the storage unit offers fault-tolerant data storage. The storage unit supports RAID implementation on the storage unit device adapters. The storage unit supports groups of disk drive modules (DDMs) in both RAID 5 and RAID 10.

RAID 5

RAID 5 is a method of spreading volume data plus data parity across multiple disk drives. RAID 5 increases performance by supporting concurrent accesses to the multiple DDMs within each logical volume.

RAID 10

RAID 10 implementation provides data mirroring from one DDM to another DDM. RAID 10 stripes data across half of the disk drives in the RAID 10 configuration. The other half of the array mirrors the first set of disk drives. In some cases, RAID 10 offers faster data reads and writes than RAID 5 because it does not need to manage parity. However, with half of the DDMs in the group used for data and the other half used to mirror that data, RAID 10 disk groups have less capacity than RAID 5 disk groups.

Storage System LPARs (logical partitions)

The DS8300 Model 9A2 exploits LPAR technology, allowing you to run two separate storage server images.

Each Storage System LPAR has access to:

- 50 percent of the processors
- 50 percent of the processor memory
- Up to 16 host adapters

- Up to 320 disk drives (up to 96 TB of capacity)

With these separate resources, each Storage System LPAR can run the same or different versions of microcode, and can be used for completely separate production, test, or other unique storage environments within this single physical system. This may enable storage consolidations where separate storage subsystems were previously required, helping to increase management efficiency and cost effectiveness.

Copy Services

Copy Services is a collection of functions that provides disaster recovery, data migration, and data duplication functions. Copy Services runs on the IBM TotalStorage DS8000 storage unit and supports open systems and zSeries environments.

Many design characteristics and advanced functions of the DS8000 contribute to protection of your data. DS8000 has a number of advanced Copy Services functions that are part of the IBM TotalStorage Resiliency family. These functions are supported also on the previous generation of storage systems called the IBM TotalStorage Enterprise Storage Server (ESS).

Copy Services include the following types of functions:

- FlashCopy, which is a point-in-time copy function
- Remote mirror and copy functions (previously known as Peer-to-Peer Remote Copy or PPRC), which includes:
 - IBM TotalStorage Metro Mirror (previously known as Synchronous PPRC)
 - IBM TotalStorage Global Copy (previously known as PPRC Extended Distance)
 - IBM TotalStorage Global Mirror (previously known as Asynchronous PPRC)
- z/OS Global Mirror (previously known as Extended Remote Copy [XRC])

You can manage Copy Services functions through a command-line interface called the IBM TotalStorage DS CLI and a new Web-based interface called the IBM TotalStorage DS Storage Manager. The DS Storage Manager allows you to set up and manage the following types of data-copy functions from any point from which network access is available:

FlashCopy

The FlashCopy feature enables you to create full volume copies of data.

When you set up a FlashCopy operation, a relationship is established between source and target volumes, and a bitmap of the source volume is created. Once this relationship and a bitmap are created, the target volume can be accessed as though all the data had been physically copied. While a relationship between the source and target volume exists, a background process copies the tracks from the source to the target volume.

FlashCopy is an optional function. To use it, you must purchase the point-in-time 2244 function authorization model, which is 2244 Model PTC.

Remote Mirror and Copy

The remote mirror and copy feature is a flexible data mirroring technology that allows replication between volumes on two or more disk storage systems. You can also use this feature for data backup and disaster recovery.

Remote mirror and copy is an optional function. To use it, you must purchase the remote mirror and copy 2244 function authorization model, which is 2244 Model RMC.

DS8000 storage units can participate in remote mirror and copy solutions with the ESS Model 750, ESS Model 800, and DS6000 storage units.

The remote mirror and copy feature can operate in the following modes:

Metro Mirror

Provides real-time mirroring of logical volumes between two DS8000s that can be located up to 300 km from each other. It is a synchronous copy solution where write operations are completed on both copies (local and remote site) before they are considered to be completed.

Global Copy

Copies data nonsynchronously and over longer distances than is possible with Metro Mirror. When operating in Global Copy mode, the source volume sends a periodic, incremental copy of updated tracks to the target volume instead of a constant stream of updates. This causes less impact to application writes for source volumes and less demand for bandwidth resources, while allowing a more flexible use of the available bandwidth.

Global Mirror

Provides a long-distance remote copy feature across two sites using asynchronous technology. Global Mirror operations provide the following benefits:

- Support for virtually unlimited distance between the local and remote sites, with the distance typically limited only by the capabilities of the network and the channel extension technology. This "unlimited" distance enables you to choose your remote site location based on business needs and enables site separation to add protection from localized disasters.
- A consistent and restartable copy of the data at the remote site, created with minimal impact to applications at the local site.
- Data currency where, for many environments, the remote site lags behind the local site on an average of 3 to 5 seconds, minimizing the amount of data exposure in the event of an unplanned outage. The actual lag in data currency that you experience can depend upon a number of factors, including specific workload characteristics and bandwidth between the local and remote sites.
- Dynamic selection of the desired recovery point objective, based upon business requirements and optimization of available bandwidth.
- Session support whereby data consistency at the remote site is internally managed across up to eight DS8000 machines that are located across the local and remote sites.
- Efficient synchronization of the local and remote sites with support for failover and failback modes, helping to reduce the time that is required to switch back to the local site after a planned or unplanned outage.

z/OS Global Mirror

The z/OS Global Mirror function mirrors data on the storage unit to a

remote location for disaster recovery. It protects data consistency across all volumes that you have defined for mirroring. The volumes can reside on several different storage units. The z/OS Global Mirror function can mirror the volumes over several thousand kilometers from the source site to the target recovery site. DS8000 storage complexes support z/OS Global Mirror only on zSeries hosts.

With z/OS Global Mirror, you can suspend or resume service during an outage. You do not have to end your current data-copy session. You can suspend the session, then restart it. Only data that changed during the outage needs to be resynchronized between the copies.

The z/OS Global Mirror function is an optional function. To use it, you must purchase the remote mirror for z/OS 2244 function authorization model, which is 2244 Model RMZ.

z/OS Metro/Global Mirror (3-site z/OS Global Mirror and Metro Mirror)

This mirroring capability uses z/OS Global Mirror to mirror primary site data to a location that is a long distance away and also uses Metro Mirror to mirror primary site data to a location within the metropolitan area. This enables a z/OS 3-site high availability and disaster recovery solution for even greater protection from unplanned outages.

The z/OS Metro/Global Mirror function is an optional function. To use it, you must purchase both of the following functions:

- Remote mirror for z/OS (2244 Model RMZ)
- Remote mirror and copy function (2244 Model RMC) for both the primary and secondary storage units

FlashCopy

The IBM TotalStorage FlashCopy feature provides a point-in-time copy capability for logical volumes. FlashCopy creates a physical point-in-time copy of the data, with minimal interruption to applications, and makes it possible to access immediately both the source and target copies.

The primary objective of FlashCopy is to create a copy of a source volume on the target volume. This copy is called a point-in-time copy. When you initiate a FlashCopy operation, a FlashCopy relationship is created between the source volume and target volume. A FlashCopy relationship is a "mapping" of a FlashCopy source volume and a FlashCopy target volume. This mapping allows a point-in-time copy of the source volume to be copied to the target volume. The FlashCopy relationship exists between the volume pair from the time that you initiate a FlashCopy operation until the DS8000 copies all data from the source volume to the target volume or until you delete the FlashCopy relationship, if it is a persistent FlashCopy.

The point-in-time copy that is created by FlashCopy is typically used when you need a copy of the production data to be produced with minimal application downtime. It can be used for online back up, testing of new applications, or for creating a database for data-mining purposes. The copy looks exactly like the original source volume and is an instantly available, binary copy.

FlashCopy supports the following copy options:

Data Set FlashCopy

Data Set FlashCopy allows a FlashCopy of a data set in a zSeries environment.

Multiple relationship FlashCopy

Multiple relationship FlashCopy allows a source to have FlashCopy relationships with multiple targets simultaneously. This flexibility allows you to establish up to 12 FlashCopy relationships on a given logical unit number (LUN), volume, or data set, without needing to first wait for or cause previous relationships to end.

Refresh target volume (also known as incremental FlashCopy)

Refresh target volume provides the ability to "refresh" a LUN or volume involved in a FlashCopy relationship. When a subsequent FlashCopy operation is initiated, only data that updates the target and the source to the same point-in-time state is copied. The direction of the "refresh" can also be reversed. The LUN or volume that was defined as the target can now become the source for the LUN or the volume that was defined as the source (now the target).

Persistent FlashCopy

Persistent FlashCopy allows the FlashCopy relationship to remain even after the FlashCopy operation completes. You must explicitly delete the relationship.

Establish FlashCopy on existing Remote Mirror and Copy source

The establish FlashCopy on an existing Remote Mirror and Copy source volume option allows you to establish a FlashCopy relationship where the target volume is also the source of an existing remote mirror and copy source volume. This enables you to create full or incremental point-in-time copies at a local site and then use remote mirroring commands to copy the data to the remote site.

This feature is represented by the **Establish target on existing Metro Mirror source** selection in the GUI.

Consistency group commands

Consistency group commands allow the DS8000 to freeze I/O activity to a LUN or volume until you issue the FlashCopy consistency group command. Consistency groups help create a consistent point-in-time copy across multiple LUNs or volumes, and even across multiple DS8000 systems. This function is available through the use of command-line interface commands.

Inband commands over remote mirror link

In a remote mirror environment, inband commands are issued to a source volume of a remote mirror and copy volume pair on a local storage unit and sent across paths (acting as a conduit) to a remote storage unit to enable a FlashCopy pair to be established at the remote site. This eliminates the need for a network connection to the remote site solely for the management of FlashCopy. This function is available through the use of command-line interface commands.

Disaster recovery using Copy Services

One of the main reasons for using Copy Services functions is in preparation for a possible disaster by backing up, copying, and mirroring your data at the production and recovery sites.

You can create a disaster recovery solution using any of the following functions. These functions help improve a disaster recovery process by allowing a much shorter recovery time with little or no data loss.

Note: In a remote mirror and copy environment, when two storage units are set up in two geographically distinct locations, we refer to the sites as Site A (the production site) and Site B (the recovery site).

Failover and failback operations

A failover is the process of temporarily switching production to a backup facility (normally your recovery site) following a planned outage, such as a scheduled maintenance period or an unplanned outage, such as a disaster. A failback operation is the process of returning production to its original location. These operations use remote mirror and copy functions to help reduce the time that is required to synchronize volumes after sites are switched during planned or unplanned outages.

Global Mirror

The Global Mirror function provides a two-site extended distance remote copy option for disaster recovery and backup solution for the zSeries and open systems environments. This solution is based on existing Global Copy and FlashCopy functions. With Global Mirror, the data that the host writes to the storage unit at the production site is asynchronously shadowed to the storage unit at the recovery site. Global Mirror operates over high-speed, Fibre Channel communication links and is designed to maintain a consistent and restartable copy of data at a recovery site that can be located at virtually unlimited distance from the production site.

Comparison of licensed functions

A key decision that you must make in planning for a disaster is deciding which licensed functions to use to best suit your environment.

Table 2 provides a brief summary of the characteristics of the Copy Services features that are available for the storage unit.

Table 2. Comparison of licensed functions

Licensed function	Description	Advantages	Considerations
Metro Mirror	Synchronous data copy at a distance	No data loss, rapid recovery time for distances up to 300 km.	Slight performance impact.
Global Copy	Continuous copy without data consistency	Nearly unlimited distance, suitable for data migration, only limited by network and channel extenders capabilities.	Copy is normally fuzzy but can be made consistent through synchronization.
Global Mirror	Asynchronous copy	Nearly unlimited distance, scalable, and low recovery point objective (RPO). The RPO is the time needed to recover from a disaster; that is, the total system downtime.	RPO might grow when link bandwidth capability is exceeded.

Table 2. Comparison of licensed functions (continued)

Licensed function	Description	Advantages	Considerations
z/OS Global Mirror	Asynchronous copy controlled by z/OS host software	Nearly unlimited distance, highly scalable, and very low RPO.	Additional host server hardware and software is required. The RPO might grow if bandwidth capability is exceeded or host performance might be impacted.

Parallel Access Volumes

The use of parallel access volumes (PAVs) enables a single zSeries server to simultaneously process multiple I/O operations to the same logical volume, which can help to significantly reduce device queue delays. This is achieved by defining multiple addresses per volume.

With dynamic parallel access volumes, the assignment of addresses to volumes is automatically managed to help the workload meet its performance objectives and minimize overall queuing.

You must configure both your storage unit and operating system to use PAVs. You can use the logical configuration definition to define PAV-bases, PAV-aliases, and their relationship in the storage unit hardware. This unit address relationship creates a single logical volume, allowing concurrent I/O operations.

The storage unit supports concurrent or parallel data transfer operations to or from the same volume from the same system or system image for S/390 or zSeries hosts. An S/390 with PAV software support enables multiple users and jobs to simultaneously access a logical volume. Read and write operations can be accessed simultaneously to different domains. (The domain of an I/O operation is the specified extents to which the I/O operation applies.)

Chapter 2. Attaching to a Apple Macintosh server

This chapter provides information about attaching an Apple Macintosh server to a storage unit.

Supported fibre-channel adapters for the Apple Macintosh server

You can use these fibre-channel adapters with the Apple Macintosh server.

You can attach a storage unit to an Apple Macintosh server with the following fibre-channel adapters:

- Apple fibre channel PCI-X card, 065-5136
- ATTO 3300
- ATTO 3321

For information about servers, operating systems, adapters, and switches that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Chapter 3. Attaching to a Fujitsu PRIMEPOWER host system

This chapter describes how to attach a Fujitsu PRIMEPOWER host to a storage unit.

Supported fibre-channel adapters for PRIMEPOWER

The following adapter card is supported for the PRIMEPOWER host system:

- Emulex LP9002L adapter card

Fibre-channel attachment requirements for PRIMEPOWER

This section lists the requirements for attaching the storage unit to your PRIMEPOWER host system.

- Ensure that there are enough fibre-channel adapters installed in the server to manage the total number of LUNs that you want to attach.
- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *User's Guide* is on the TotalStorage CD that you receive with the storage unit.
- Review device driver installation documents and configuration utility documents for any PRIMEPOWER patches that you might need.
- See the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html> for details about the release level for your operating system.

Either you or an IBM service support representative must perform the following tasks to install and configure a storage unit:

1. Install the storage unit.
2. Define the fibre-channel host system with the worldwide port name identifiers. For the list of worldwide port names, see "Locating the worldwide port name (WWPN)," on page 225.
3. Define the fibre-port configuration if you did not do it when you installed the storage unit or fibre-channel adapters.
4. Configure the host system for the storage unit by using the instructions in your host system publications.

Installing the Emulex adapter card for a PRIMEPOWER host system

This section tells you how to attach a storage unit to a PRIMEPOWER host system with an Emulex adapter card.

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

- Target
 - Public initiator
 - Private initiator
 - Target and public initiator
 - Target and private initiator
1. Record the IEEE number that is printed on the card. You can use the IEEE number to determine the WWPN.

2. Refer to the installation instructions provided by your host adapter vendor for their specific installation instructions. See <http://www.emulex.com> for the latest documentation.

Downloading the Emulex adapter driver for a PRIMEPOWER host system

This section provides instructions to download the current Emulex fibre-channel adapter driver.

1. Restart your host system.
2. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
3. Click **Interoperability matrix**.
4. Click **DS8000 interoperability matrix**.
5. Find the section for the current version of the driver and firmware that you want.
6. Go to <http://www.emulex.com>.
7. Click **Driver, downloads, and documentation** from the left navigation pane. Click **OEM software and documentation**.
8. Click **IBM**.
9. Click the link for the adapter that corresponds to the firmware, driver, and documentation that you need to install and download the adapter driver.

Installing the Emulex adapter driver for a PRIMEPOWER host system

This section provides instructions to install the Emulex fibre-channel adapter driver.

1. Go to <http://www.emulex.com>.
2. From the left navigation pane, click **Drivers, Software, and Manuals**.
3. Click **IBM**.
4. Determine the driver you want, and click the link for the adapter driver.
5. Determine the operating system for the driver that you want to download and click **Version x-x.xxxx**, where x-x.xxxx is the version number for the adapter driver.

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

Configuring host device drivers for PRIMEPOWER

Perform the following steps to update the PRIMEPOWER driver configuration file. This procedure gives you access to target and LUN pairs that are configured on the storage unit.

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

1. Change to the directory by typing: `cd /kernel/drv`
2. Backup the `sd.conf` file in this subdirectory.
3. Edit the `sd.conf` file to add support for the target and LUN pairs that are configured on the host system.

Note: Do not add duplicate target and LUN pairs.

Figure 9 shows the lines that you must add to the file to access LUNs 0 - 49 on target 0 for fibre-channel.

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

Figure 9. Example of sd.conf file entries for fibre-channel

Figure 10 on page 34 shows the start lpfc auto-generated configuration.

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

```

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

```

Figure 10. Example of a start lpfc auto-generated configuration

4. Type either:

- a. `reboot -- -r` from the Open Windows window to shutdown and restart the PRIMEPOWER host system with the kernel reconfiguration option
- b. `boot -r` from the OK prompt after you shutdown

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

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

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

Parameter settings for the Emulex LP9002L adapter

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

Table 3 on page 35 provides configuration settings that are recommended for the Emulex LP9002L adapter. For the most current information about fibre-channel adapter parameter settings, see:

<http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do>

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

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

Setting parameters for Emulex adapters

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

1. Type `cd /etc` to change to the `/etc` subdirectory.
2. Backup the system file in the subdirectory.
3. Edit the system file and set the following parameters for servers with configurations that only use Emulex or QLogic adapters.

sd_io_time

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

sd_retry_count

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

maxphys

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

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

Chapter 4. Attaching to a Hewlett-Packard AlphaServer Tru64 UNIX host

This chapter describes how to attach a Hewlett-Packard (HP) AlphaServer Tru64 UNIX host to a storage unit.

Attaching to an HP AlphaServer Tru64 UNIX host system with fibre-channel adapters

This section describes the host system requirements and provides the procedures to attach a storage unit to an HP AlphaServer with fibre-channel adapters.

Supported fibre-channel adapters for the HP AlphaServer Tru64 UNIX host system

This section lists the fibre-channel adapters that are supported for storage unit attachment to an HP AlphaServer Tru64 UNIX host system.

You can attach a storage unit to an HP AlphaServer Tru64 UNIX host system with the following fibre-channel adapters:

- KGPSA-CA adapter card
- KGPSA-DA adapter card
- KGPSA-EA adapter card

Note: You do not need the Subsystem Device Driver because Tru64 UNIX manages multipathing.

For information about adapters that you can use to attach the storage unit to the HP host and the HP AlphaServer models, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Supported operating system levels for fibre-channel attachment to an HP Tru64 UNIX host

This section lists the supported operating system levels for storage unit fibre-channel attachment to an HP AlphaServer Tru64 UNIX host system.

You can attach a clustered HP AlphaServer Tru64 host system with level 5.0A, 5.1, 5.1A, and 5.1B with the following topologies:

- Switched fabric
- Point-to-point

Fibre-channel Tru64 UNIX attachment requirements

This section lists the attachment requirements to attach a storage unit to a Tru64 UNIX host system with fibre-channel adapters.

You must comply with the following requirements when attaching the storage unit to your host system:

- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *User's Guide* is on the TotalStorage CD that you receive with the storage unit.

- See the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html> for details about the release level for your operating system.
- An IBM service support representative (SSR) must install the storage unit.

Fibre-channel Tru64 UNIX attachment considerations

This section lists the attachment considerations when attaching a storage unit to an HP AlphaServer Tru64 UNIX host system.

Table 4 lists the maximum number of adapters you can have for an AlphaServer.

Note: For a list of open systems hosts, operating systems, adapters and switches that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Table 4. Maximum number of adapters you can use for an AlphaServer

AlphaServer name	Maximum number of adapters
800	2
1200	4
2100	4
4000, 4000a	4
4100	4
8200, 8400	8
DS10, DS20, DS20E	2
ES40, ES45, ES47	4
GS60, GS60E, GS140	8
GS80, GS160, GS320	8

Supporting the AlphaServer Console for Tru64 UNIX

The following sections describe how to support the AlphaServer Console for Tru64 UNIX when you are attaching to a storage unit with fibre-channel adapters.

Supported microcode levels for the HP Tru64 UNIX host system

This topic lists the supported storage unit microcode levels when you attach to an HP AlphaServer Tru64 UNIX host system.

Support for the AlphaServer console that recognizes the storage unit LUNs with fibre is available for the current version of the storage unit that is LMC. To determine which level of LMC is running on the storage unit, type `telnet xxxxx`, where `xxxxx` is the storage unit cluster name.

Figure 11 on page 39 shows an example of the output from the **telnet** command.

```

telnet xxxxx (where xxxxx is the cluster name)

      IBM TotalStorage DS8000

Model 921          SN 75-99999      Cluster Enclosure 1

OS Level 4.3.2.15  Code EC 1.5.0.107  EC Installed on: Jan 10 2005

      tbld1212
      SEA.rte level = 2.6.402.678
      SEA.ras level = 2.6.402.678

      licensed machine code - Property of IBM.

      2107 licensed machine code
      (C) IBM Corporation 1997, 2004. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure

      Restricted by GSA ADP Schedule Contract with IBM Corporation.

Login:

```

Figure 11. Confirming the storage unit licensed machine code on an HP AlphaServer through the telnet command

Supported switches for the HP Tru64 UNIX host system

This section lists the switches that are supported when you attach a storage unit to an HP AlphaServer Tru64 UNIX host system.

For more information on supported switches, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>

IBM supports cascaded switches in configurations up to a maximum of 8 switches with a maximum of 3 interswitch hops for any path. Use two hops for normal operations with the third hop reserved for backup paths.

Attaching the HP AlphaServer Tru64 UNIX host to a storage unit using fibre-channel adapters

To attach a storage unit to a Tru64 UNIX operating system, perform the following tasks:

- Confirm the installation of the operating system
- Confirm the installation of the host adapter
- Set up the storage unit
- Confirm switch connectivity
- Confirm storage connectivity
- Configure storage
- Review tuning recommendations
- Review unsupported utilities

Confirming the installation of the Tru64 UNIX operating system

This section lists the steps you must perform to confirm the installation of Tru64 UNIX.

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

1. Confirm the installation of the appropriate version of Tru64 UNIX. For the Tru64 UNIX 5.x operating system, use the **sizer -v** command to confirm installation. Figure 12 shows an example of what is displayed when you use the **sizer -v** command.

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

Figure 12. Example of the sizer -v command

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

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

This section lists the steps you must perform to install the KGPSA-xx adapter card.

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

See Figure 13 on page 41 for an example of what you see when you type `set mode diag` and `wwidmgr -show adapter`.


```

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

```

Figure 13. Example of the set mode diag command and the wwidmgr -show adapter command

Figure 13 shows the worldwide node name (WWNN). You need the worldwide port name (WWPN) to configure the storage unit host attachment. To determine the WWPN for the KGPSA adapters, replace the “2” in the WWNN with a “1”.

Setting the mode for the KGPSA-xx host adapter

This task describes setting the mode for the KGPSA-xx host adapter.

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

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

1. Type # shutdown -h now to shutdown the operating system.
2. Type init to initialize the system.
3. Type wwidmgr -show adapter to check the mode. Figure 14 shows example output from the **wwidmgr** command.

```

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

```

Figure 14. Example results of the wwidmgr command.

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

Setting up the storage unit to attach to an HP AlphaServer Tru64 UNIX host system with fibre-channel adapters

The following sections tell you how to set up a storage unit to attach to an HP AlphaServer Tru64 UNIX host.

Adding or modifying AlphaServer fibre-channel connections

This section provides the reference you need to add or modify AlphaServer fibre-channel connections.

To add, remove, or modify the AlphaServer connections, use the IBM TotalStorage DS Storage Manager. When you add a connection, it is necessary to specify the worldwide port name of the host connection. See “Locating the worldwide port name (WWPN),” on page 225 for procedures on how to locate the WWPN for each KGPSA adapter card.

Configuring fibre-channel host adapter ports for the Tru64 UNIX host system

This section provides the reference you need to configure host adapter ports for fibre-channel storage unit connectivity to an HP AlphaServer Tru64 UNIX host system.

To configure the host adapter ports, use the IBM TotalStorage DS Storage Manager or see the *IBM TotalStorage DS8000 Command-Line Interface User's Guide*.

Adding and assigning volumes to the Tru64 UNIX host system

This section provides information you need to add and assign volumes to the Tru64 UNIX host system.

To set up disk groups, create volumes and assign them to the Tru64 connections, use the IBM TotalStorage DS Storage Manager or see the *IBM TotalStorage DS8000 Command-Line Interface User's Guide*.

Confirming fibre-channel switch connectivity for Tru64 UNIX

This section lists the steps you need to perform to confirm fibre-channel switch connectivity for Tru64 UNIX.

1. Open a telnet session and log in to the switch as an administrator.
2. Confirm that each host adapter has performed a fabric login to the switch.
3. Type `switchshow` to confirm that each storage unit host adapter has performed a fabric login to the switch. Figure 15 on page 43 shows an example of what displays when you type the **switchshow** command.

```

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

```

Figure 15. Example of the switchshow command

Confirming fibre-channel storage connectivity for Tru64 UNIX

This section lists the steps you must perform to confirm fibre-channel storage connectivity for an HP Tru64 UNIX host system.

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

Type wwidmgr -show wwid to display the information about the storage unit volume at the console level. You can use this information to identify the volumes that are attached to an AlphaServer. Figure 16 shows an example of information about the storage unit volumes that you can see at the AlphaServer console.

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

```

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>>>

```

Figure 16. Example of storage unit volumes on the AlphaServer console

See “Tru64 UNIX UDID hexadecimal representations” on page 44 for an explanation of the UDID.

Tru64 UNIX UDID hexadecimal representations

This section describes UDID representations for storage unit volumes.

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

The underscore in Figure 17 highlights the hex string that identifies a storage unit volume that is attached to an AlphaServer.

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

Figure 17. Example of a hex string for a storage unit volume on an AlphaServer Tru64 UNIX console

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

The underscore in Figure 18 highlights an example of a hex string that identifies the decimal volume number of the storage unit volume. The first three characters of the next to last quartet of numbers is the hex string representation. Figure 18 shows that the storage unit volume number is decimal 282.

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

Figure 18. Example of a hex string that identifies the decimal volume number for a storage unit volume on an AlphaServer console or Tru64 UNIX

Figure 19 shows a hex representation of the last 5 characters of the storage unit volume serial number.

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

Figure 19. Example of hex representation of last 5 characters of a storage unit volume serial number on an AlphaServer console

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

This section lists the steps you must perform to prepare to boot from the storage unit.

Use the **wwidmgr** command to set up each device that you use for booting or dumping. After you set up a device, the console retains the information that is needed to access the device in nonvolatile memory. Rerun the **wwidmgr** command if the system configuration changes and the nonvolatile information is no longer valid.

1. Display the WWIDs of all assigned storage unit volumes with the **wwidmgr -show wwid** command.

- Determine which storage unit volume that you want to use as a boot or dump device by decoding the serial number as described in “Tru64 UNIX UDID hexadecimal representations” on page 44.
- Assign a unit number with the **wwidmgr -quickset -item *i* -unit *u*** command, where *i* is the wwidmgr item number and *u* is the unit number you choose. You can find the item number inside the square brackets of the output from the **wwidmgr -show wwid** command.
- Reinitialize the server with the **init** command.
When you make changes with the **wwidmgr** command, they do not take effect until the next system initialization.
- Use the **show device** command to verify that the system displays the disk as console device **DGAu**, with the unit number that you defined.
After the initialization, the console **show device** command displays each fibre-channel adapter followed by the paths through that adapter to each of the defined fibre-channel disks. The path-independent OpenVMS device name for each fibre-channel disk is displayed in the second column.
- Set the default boot device console variable, **bootdef_dev**, to match the console device name of the boot disk. In a multipath environment, use the asterisk (*) as a wildcard to make all boot paths available. An example of the multipath command is **set bootdef_dev DGAu.***, where *u* is the unit number and * denotes all possible paths.

Verifying the fibre-channel attachment of the storage unit volumes to Tru64 UNIX

This section contains steps you must perform to verify the fibre-channel attachment of the storage unit volumes to the Tru64 UNIX host system.

For Tru64 UNIX 5.x:

- Use the **hwmgr** command to verify the attachment of the storage unit volumes for Tru64 5.x. Figure 20 shows an example of the commands you can use to verify the attachment of the storage unit volumes.

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

Figure 20. Example of the hwmgr command to verify attachment

2. Use the example Korn shell script, called `dsvol`, shown in Figure 21, to display a summary that includes information for all the storage unit volumes that are attached.

```
echo Extracting DS volume information...
for ID in `hwmgr -view dev -cat disk | grep ibm2107 | awk '{ print $1}'`
do echo; echo DS vol, H/W ID $ID
hwmgr -get attrib -id $ID | awk '/phys_loc//dev_base//capacity//serial/'
done
```

Figure 21. Example of a Korn shell script to display a summary of storage unit volumes

Figure 22 shows an example of what displays when you execute the `dsvol` Korn shell script.

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

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

```
# ./dsvol | more
Extracting DS volume information...ESS vol, H/W ID 38:
phys_location = bus-2-targ-0-lun-0
dev_base_name = dsk3
capacity = 5859392
serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-2821-5660
DS vol, H/W ID 39:
phys_location = bus-2-targ-0-lun-1
dev_base_name = dsk4
capacity = 5859392
serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-2831-5660
DS vol, H/W ID 40:
phys_location = bus-2-targ-0-lun-2
dev_base_name = dsk5
capacity = 5859392
serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-2841-5660
#
```

Figure 22. Example of the Korn shell script output

Configuring the storage for fibre-channel Tru64 UNIX hosts

This section lists the steps you must perform to configure the storage for fibre-channel Tru64 UNIX host systems.

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

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

1. Perform the following steps to configure an AdvFS file system:
 - a. Type: `# disklabel -wr /dev/rdisk/dsk6c`
 - b. Type: `# mkfdmn /dev/disk/dsk6c adomain`

- c. Type: # mkfset adomain afs
 - d. Type: # mkdir /fs
 - e. Type: # mount -t advfs adomain#afs /fs
2. Perform the following steps to configure an Ufs file system:
- a. Type: # disklevel -wr /dev/disk/dsk6c
 - b. Type: # newfs /dev/disk/dsk6c
 - c. Type: # mkdir /fs
 - d. Type: # mount -t ufs /dev/disk/dsk6c /fs

Removing persistent reserves for Tru64 UNIX 5.x

This section explains how to remove persistent reserves for the Tru64 UNIX 5.x host system.

In the clustered environment, Tru64 5.x hosts place a persistent reserve whenever you assign a LUN. If you perform a FlashCopy or Remote Mirror and Copy operation on a LUN that has persistent reserve, it will fail. If you want to perform a FlashCopy or Remote Mirror and Copy operation, remove the persistent reserve on the target LUN before you assign the LUN.

For example, assume that there are two Tru64 5.x hosts, Alpha1 and Alpha2. The connections available for accessing storage unit LUNs with fibre:

- **Fibre connection:** One fibre connection goes from the KGPSA-xx card in Alpha1 to a switch. One fibre connection goes from the Alpha2 to the switch and another fibre connection goes from the switch to the storage unit.

Use a storage unit volume, for example 10a-21380, as a target LUN to perform a FlashCopy. There are two hosts and four connections: two from Alpha1 and two from Alpha2. The storage unit volume 10a-21380 has four registrants. One registrant is reserved. Use the essvol script to find the devices that are associated with storage unit volume 10a-21380 on each Trucluster node. Figure 23 on page 48 shows an example of how to remove a persistent reserve when you use the essvol script.

```

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

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

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

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

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

```

Figure 23. Example of the *essvol* script

Use the **scu** command to see the reservations on these devices. Figure 24 and Figure 25 on page 49 show examples of what you see when you use the **scu** command. You can associate dsk47 on Alpha1 and Alpha2 with storage unit volume 10a-21380.

```

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

      Generation Value: 49
      Additional Length: 16

Reservation Descriptors:

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

```

Figure 24. Example of the *scu* command


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

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

```
Reservation Descriptors:
```

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

Figure 25. Example of the scu command

Each device shows 0x30001 as a reservation key. You must issue the **scu** command again to remove the persistent reserve using reservation key 0x30001 on each node. Tru64 places a unique reservation key on each LUN whenever the storage unit assigns the LUN. The reservation key can only be removed from that specific host to which it is assigned. Because it is not possible to tell exactly which registrant on the host holds the reservation, you must issue an **scu clear** command on each node in the cluster. Figure 26 shows an example of what you see when you use the **scu clear** command.

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

Figure 26. Example of the scu clear command

One of the two commands that you see in Figure 26 clears the persistent reserve on storage unit volume 10a-21380.

Use the **scu** command to check reservations. Figure 27 on page 50 shows an example of what you see when you use the **scu** command again.

One of the two commands you see in Figure 27 on page 50 clears the persistent reserves on storage unit volume 10a-21380.

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

Persistent Reservation Header:

Generation Value: 50

Additional Length: 0

Reservation Descriptors:

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

Persistent Reservation Header:

Generation Value: 50

Additional Length: 0

Reservation Descriptors:

Figure 27. Example of the scu command to show persistent reserves

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

Limitations for Tru64 UNIX

The following is a list of limitations for Tru64 UNIX for fibre-channel connections:

Boot Volumes

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

UFS file system

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

AdvFS file system

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

From the command line, type the following commands:

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

Chapter 5. Attaching to a Hewlett-Packard AlphaServer OpenVMS host

This chapter describes how to attach a Hewlett-Packard (HP) AlphaServer OpenVMS host to a storage unit with fibre-channel adapters.

Supported fibre-channel adapters for the HP AlphaServer OpenVMS host system

You can attach a storage unit to an HP AlphaServer that is running the OpenVMS operating system with the following fibre-channel adapters:

- KGPSA-CA adapter card
- KGPSA-DA (FCA-2354) adapter card
- KGPSA-EA (FCA-2384) adapter card

Note: You do not need the Subsystem Device Driver because OpenVMS manages multipathing.

For information about adapters that you can use to attach the storage unit to the HP host and the HP AlphaServer models, see <http://www-1.ibm.com/servers/storage/disk/ds8000>.

Fibre-channel OpenVMS attachment requirements

Review the following requirements before you attach the storage unit to your host system:

- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *User's Guide* is on the TotalStorage CD that you receive with the storage unit.
- See the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html> for details about the release level for your operating system.
- An IBM service support representative must install and configure an IBM storage unit.

Fibre-channel OpenVMS attachment considerations

Table 5 lists the maximum number of adapters you can have for an AlphaServer.

Table 5. Maximum number of adapters you can use for an AlphaServer

AlphaServer name	Maximum number of adapters
800	2
1200	4
4000, 4000a	4
4100	4
8200, 8400	8
DS10, DS20, DS20E	2
ES40, ES45, ES47	4

Table 5. Maximum number of adapters you can use for an AlphaServer (continued)

AlphaServer name	Maximum number of adapters
GS60, GS60E, GS140	8
GS80, GS160, GS320	8

Supported OpenVMS feature codes

For more information about the feature codes for the storage unit and the distances that are supported by fibre-channel cables for the storage unit, see *IBM TotalStorage DS8000 Introduction and Planning Guide*.

Supported microcode levels for the HP OpenVMS host system

For more information about the microcode levels for the storage unit, see *IBM TotalStorage DS8000 Introduction and Planning Guide*.

Supported switches for the HP OpenVMS host system

IBM supports cascaded switches in configurations up to a maximum of 8 switches with a maximum of 3 interswitch hops for any path. Use two hops for normal operations with the third hop reserved for backup paths.

For more information on supported switches, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>

Attaching the HP AlphaServer OpenVMS host to a storage unit using fibre-channel adapters

The following sections describe attaching the HP AlphaServer OpenVMS host to a storage unit using fibre-channel adapters:

- Confirm the installation of the operating system
- Installing the KGPSA-xx adapter card
- Set up the storage unit
- Confirm switch connectivity
- Confirm storage connectivity
- Configure storage
- Restrictions
- Troubleshooting fibre-attached volumes

Supported operating system levels for fibre-channel attachment to an HP OpenVMS host

This section lists the supported operating system levels for fibre-channel attachment to an OpenVMS host system.

You can attach a storage unit to an OpenVMS host with the following system levels:

- 7.3
- 7.3-1
- 7.3-2

See the HP document, *Guidelines for OpenVMS Cluster Configurations*, for a discussion about working with fibre-channel devices and system parameters such as MSCP_CMD_TMO, MVTIMEOUT, and MPDEV_LCRETRIES.

Confirming the installation of the OpenVMS operating system

This section provides steps to confirm the installation of the OpenVMS operating system.

Ensure that the storage unit supports attachment to the OpenVMS operating system level that you are installing. See “Supported operating system levels for fibre-channel attachment to an HP OpenVMS host” on page 52 for a list of supported OpenVMS operating system levels.

1. Use the **show system** command to show the current version of the OpenVMS operating system that you have installed. Figure 28 shows an example of what is displayed when you use the **show system** command.

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

Figure 28. Example of show system command to show system command on the OpenVMS operating system

2. Ensure that you have installed the most current version of patches and HP-recommended remedial kits. Some kits are dependent on other kits. You must install some kits before you install another kit. See the documentation from HP for more information.
3. Use the **product show history** command to check the patches that are installed on your system. Figure 29 shows an example of what you see when you use the **product show history** command.

```
$ product show history
```

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

```
8 items found
$
```

Figure 29. Example of the product show history command to check the versions of patches already installed

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

This section provides the steps you must follow to install the KGPSA-xx adapter card.

1. Shut down the Hewlett-Packard AlphaServer host system.

2. Use the procedures that are provided by the manufacturer of the adapter card to install the KGPSA-xx host adapter.
3. Restart the host (nonclustered configurations) or each cluster member (clustered configurations).
4. Bring each host system to a halt condition at the console level.
5. Type `set mode diag` at the Hewlett-Packard AlphaServer console to place the console in diagnostic mode.

Note: This step is only for certain AlphaServer models.

6. Type `wwidmgr-show adapter` to confirm that you installed each adapter properly.
7. If necessary, update the adapter firmware.

Figure 30 shows an example of what you see when you type `set mode diag` and `wwidmgr -show adapter`. Figure 30 shows the worldwide node name (WWNN). You also need the worldwide port name (WWPN) to configure the storage unit host attachment. To determine the WWPN for the KGPSA adapters, replace the "2" in the WWNN with a "1".

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

Figure 30. Example of the `set mode diag` command and the `wwidmgr -show adapter` command

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

This task describes setting the mode for the KGPSA-xx host adapter for an HP OpenVMS host system.

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

The KGPSA-xx fibre-channel adapter card must be set to FABRIC. Use the **wwidmgr -show adapter** console command to see the setting for the Cur. Topo and Next Topo variables. Figure 31 on page 55 shows example output from the **wwidmgr** command.

1. Type `# shutdown -h now` to shutdown the operating system.
2. Place the AlphaServer into console mode.
3. Type `wwidmgr -show adapter` to check the mode. Figure 31 on page 55 shows example output from the **wwidmgr** command.

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

Figure 31. Example results of the `wwidmgr` command.

4. Type `wwidmgr -set adapter -item 9999 -topo fabric` to set the KGPSA host adapter mode to FABRIC.
5. Type `init` to initialize the system.
6. Type `wwidmgr -show adapter` to check the mode.

Setting up the storage unit to attach to an HP AlphaServer OpenVMS host system with fibre-channel adapters

This section describes setting up the storage unit to attach to an HP AlphaServer OpenVMS host system with fibre-channel adapters.

Adding or modifying AlphaServer fibre-channel connections for the OpenVMS host

To create, remove, or modify the AlphaServer connections, use the IBM TotalStorage DS Storage Manager. When you create a connection, it is necessary to specify the worldwide port name of the host connection. See “Locating the worldwide port name (WWPN),” on page 225 for procedures on how to locate the WWPN.

Defining OpenVMS fibre-channel adapters to the storage unit

This section lists the steps you must perform to define OpenVMS fibre-channel adapters to the storage unit.

You must create a storage complex before defining each host adapter to the storage unit. Use the IBM TotalStorage DS Storage Manager to create a storage complex. You can use the following menu path: click **Real-time manager**, click **Manage hardware**, click **Storage complexes**.

To define each host adapter to the storage unit, perform the following steps using the IBM TotalStorage DS Storage Manager.

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

1. Logon to the IBM TotalStorage DS Storage Manager and click **Real-time manager**, click **Manage hardware**, click **Host systems**.
2. From the **Select action box**, click **Create**, then click **Go**.
3. From the **General host information** screen, complete the following fields for each fibre-channel host adapter. When finished, click **OK**.
 - Host type
 - Nickname
 - Description
4. From the **Define host ports** screen, specify the host ports for this host. Click **Add** to add each host port to the defined host ports table.

5. From the **Define host WWPN** screen, specify the world-wide port names for the selected hosts. When finished, click **Next**.
6. From the **Select storage images** screen, specify the storage image for attachment. Click **Add** to add the storage image to the selected storage images table. When finished, click **Next**.
7. From the **Specify storage image parameters** screen, specify the following parameters for each storage image.
 - Host attachment identifier
 - Volume Group for host attachment
 - Type of I/O port (any valid storage image I/O port or a specific I/O port)
8. Click the **Apply assignment** button to apply the current attachment assignment. Use this button to go through this page repeatedly for each host attachment identifier that you want to assign to the storage image. When finished, click **OK**.
9. From the **Verification** screen, verify the data that represents the new host attachments. When finished, click **Finish**.

Configuring fibre-channel host adapter ports for OpenVMS

To configure the host adapter ports, use the IBM TotalStorage DS Storage Manager.

OpenVMS fibre-channel considerations

When you assign an OpenVMS host to a specific storage unit adapter port, ensure that you do not assign any other host type to the same storage unit adapter. Do not allow any other host that is not assigned to that adapter to access the OpenVMS-specific storage unit adapter port.

Using fabric zoning, you must create a zone strictly for OpenVMS hosts. The storage unit fibre-channel adapter port to which you assigned a storage unit host must reside exclusively within this OpenVMS zone. Failure to ensure this exclusivity might cause an OpenVMS host to hang. If this occurs, you must reboot the OpenVMS host.

OpenVMS UDID Support

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

You can use the IBM TotalStorage DS Storage Manager or the DS8000 command line interface (DS CLI) to set a value in a storage unit volume name field that is used by AlphaServer systems as the UDID for the volume. (In this document, we provide DS CLI examples.) You can find the DS CLI on the CD that you receive with the storage unit. See the *IBM TotalStorage DS8000 Command-Line Interface User's Guide* for more information.

The DS CLI is a general purpose utility that supports various storage unit functions. The DS CLI allows 16 alphanumeric characters as input when you complete the storage unit volume name field. OpenVMS UDID values must be an integer within

the range of 0 to 32767. Therefore, you must ensure that the input is valid for UDID support. The utility does not enforce UDID rules. It accepts values, such as AaBbCcDd, that are not valid for OpenVMS. It is possible to assign the same UDID value to multiple storage unit volumes. However, each volume that you assign to an OpenVMS system must have a value that is unique for that system or throughout the OpenVMS cluster to which that system belongs. Review the HP OpenVMS documentation for UDID rules, and verify that your input is valid.

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

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

1. Use the chfbvol command to change the name of a fixed block volume: For example to set the UIDD value to 21, type: `chfbvol -name 21 0001`
The value for the name parameter in the DS CLI command is the UDID field for the HP AlphaServer.

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

2. To make a volume group called "VMS_A0" and add a volume to it, type:
`mkvolgrp -type scsimap256 -volume 0001 VMS_A0`

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

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

3. To create an OpenVMS host with the DS CLI and associate a volume group with it, type: `mkhostconnect -wwname 10000000ABCDEF98 -hosttype HpVms -volgrp v0 ES40_A`

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

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

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

The following is example output:

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

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

The following is example output:

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

6. Use the `wwidmgr -show wwid` command at the AlphaServer console to list fibre-attached volumes that have been detected by its fibre-channel host adapters. If a volume has no UDID or has an invalid UDID, the volume UDID number is minus one (-1) or zero (0). When it is booted, OpenVMS does not detect a volume with -1 or 0 as a UDID number.

OpenVMS LUN 0 - Command Control LUN

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

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

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

At the AlphaServer console, when you issue the `wwidmgr -show wwid` command, LUN 0 does not display. Only fibre-channel storage devices are listed in the output. The storage unit LUN 0 is presented as a CCL device, and therefore, is not shown.

Confirming fibre-channel switch connectivity for OpenVMS

This section lists the steps you must perform to confirm fibre-channel switch connectivity for the OpenVMS host system.

1. Open a telnet session and log in to the switch as an administrator.
2. Confirm that each host adapter has performed a fabric login to the switch. Figure 32 on page 59 shows an example of what displays when you type the **switchshow** command.
3. Confirm that each storage unit host adapter has performed a fabric login to the switch. Figure 32 on page 59 shows an example of what displays when you type the **switchshow** command.

```

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

```

Figure 32. Example of the switchshow command

Confirming fibre-channel storage connectivity for OpenVMS

Perform the following steps to confirm the fibre-channel storage connectivity for the OpenVMS host system.

1. Reset the host (nonclustered configurations) or each cluster member (clustered configurations).
2. Bring each host system to a halt condition at the console level.
3. If required by the host, type set mode diag at the Hewlett-Packard AlphaServer console to place the console in diagnostic mode.

Type wwidmgr -show wwid to display the information about the storage unit volume at the console level. You can use this information to identify the volumes that are attached to an AlphaServer. Figure 33 shows an example of information about the storage unit volumes that you can see at the AlphaServer console.

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

```

P00>>>set mode diag
Console is in diagnostic mode
P00>>>wwidmgr -show wwid
[0] UDID:20 WWID:01000010:6005-0763-03ff-c0a4-0000-0000-0000-000f (ev:none)
[1] UDID:21 WWID:01000010:6005-0763-03ff-c0a4-0000-0000-0000-0001 (ev:none)
P00>>>

```

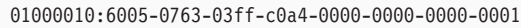
Figure 33. Example of storage unit volumes on the AlphaServer console

OpenVMS World Wide Node Name hexadecimal representations

This section explains OpenVMS World Wide Node Name hexadecimal representations.

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

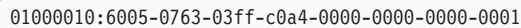
The underscore in Figure 34 highlights the World Wide Node Name that identifies the storage unit volume that is attached to an AlphaServer.



01000010:6005-0763-03ff-c0a4-0000-0000-0000-0001

Figure 34. Example of a World Wide Node Name for the storage unit volume on an AlphaServer console

The underscore in Figure 35 highlights the hex string that identifies the volume number of the storage unit volume.



01000010:6005-0763-03ff-c0a4-0000-0000-0000-0001

Figure 35. Example of a volume number for the storage unit volume on an AlphaServer console

Verifying the fibre-channel attachment of the storage unit volumes for OpenVMS

Follow these steps to verify the fibre-channel attachment of the storage unit volumes for OpenVMS.

1. Start the operating system.
2. Use the standard OpenVMS storage configuration utilities to prepare the storage unit LUNs and create and mount devices. Figure 36 on page 61 shows an example of what is displayed when you use the standard OpenVMS storage configuration utilities.

```

$ SHOW DEVICES DG

Device          Device Error Volume Free  Trans Mnt
Name            Status Count Label  Blocks Count Cnt
1$DGA20: (HANK) Online  0
1$DGA21: (HANK) Online  0
$
$ INITIALIZE/SYSTEM $1$DGA20 ESS001

$ MOUNT/SYSTEM $1$DGA20 ESS001

$ DIRECTORY $1$DGA20:[000000]

$ DISMOUNT $1$DGA20

```

Figure 36. Example of what is displayed when you use OpenVMS storage configuration utilities

Configuring the storage for fibre-channel OpenVMS hosts

Perform the following steps to configure the storage for fibre-channel OpenVMS hosts.

1. Start the operating system.
2. Initialize the storage unit volumes.

On the OpenVMS platform, you can initialize storage unit volumes as ODS-2 or ODS-5 volumes. You can use the volumes to create volume sets. Volume sets are concatenated volumes that form a larger volume. See the HP document, *OpenVMS System Manager's Manual, Volume 1: Essentials*.

3. Mount the storage unit volumes.

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

4. Access the storage unit volumes.

OpenVMS fibre-channel restrictions

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

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

- Multiple OpenVMS systems can access the same storage unit adapter port. However, you must define each system for that specific storage unit port and assign at least one storage unit volume.
- To re-enable compatibility, you can force the defined OpenVMS host to reestablish connectivity to the storage unit adapter by disabling the switch port it is connected to, and then enable the switch port again.

Troubleshooting fibre-channel attached volumes for the OpenVMS host system

This section contains a few examples of issues that can occur with your OpenVMS host system and some suggestions to correct the issues.

Problem

The following communication problems might occur when you attach a storage unit to an HP AlphaServer OpenVMS host system.

Investigation

If the system hangs during boot

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

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

If the system reaches the mount verification timeout

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

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

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

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

```
-SYSTEM-F-TOOMANYRED, too many redirects
```

Force the host adapter to reestablish connectivity

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

storage unit adapter port. You can also reestablish connectivity by accessing the fibre-channel switch and disabling one of the switch ports, and then enabling it again.

Chapter 6. Attaching to a Hewlett-Packard Servers (HP-UX) host

This chapter provides instructions, requirements, and considerations for attaching a Hewlett-Packard Servers (HP-UX) host system to a storage unit.

Attaching with fibre-channel adapters

This section describes the host system requirements and provides procedures to attach a storage unit to a Hewlett-Packard 9000 host system with fibre-channel adapters.

Supported fibre-channel adapters for HP-UX hosts

This section lists the supported fibre-channel adapters for HP-UX host systems.

This section describes how to attach a storage unit to a Hewlett-Packard host system with the following fibre-channel adapter cards:

- A5158A
- A6684A
- A6685A
- A6795A
- A6826A
- A9782A

Fibre-channel attachment requirements for HP-UX hosts

This section lists the requirements for attaching the storage unit to your HP-UX host system:

- Check the LUN limitations for your host system.
- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *User's Guide* is on the TotalStorage CD that you receive with the storage unit.
- See the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html> for details about the release level for your operating system, open-systems hosts, adapters and switches that IBM supports.

Either you or an IBM service support representative (SSR) must perform the following tasks to install and configure a storage unit:

1. Install the storage unit.
2. Define the fibre-channel host system with the worldwide port name identifiers. For the list of worldwide port names, see "Locating the worldwide port name (WWPN)," on page 225.
3. Define the fibre-port configuration if you did not do it when you installed the storage unit or fibre-channel adapters.
4. Configure the host system for the storage unit by using the instructions in your host system publications.

Note: To have failover protection on an open system, SDD requires at least 2 paths, and SDD allows a maximum of 32 paths.

Installing the fibre-channel adapter drivers for HP-UX 11.i, and HP-UX 11iv2

This section tells you how to download and configure the following fibre-channel adapter drivers for HP-UX 11.i and HP-UX 11iv2:

- A5158A
- A6684A
- A6685A
- A6975A
- A9782A

1. Go to:

<http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do>

2. Select the appropriate options for your product and operating system.
3. Find the section for the current version of the driver and firmware and driver you want.
4. Click on **View Details**.
5. Under *FC HBA drivers download*, click on the driver type that you need for your system.
6. Find the driver that you need for your system and click **Download**.
7. Follow the system instructions to install the driver.

Setting the queue depth for the HP-UX operating system with fibre-channel adapters

Before you set the queue depth, you must connect the host system to the storage unit. See “General information about attaching to open-systems host with fibre-channel adapters” on page 16.

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

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

2. You must monitor configurations with greater than 256 LUNs. You must adjust the queue depth for optimum performance.
3. To update the queue depth by device level, use the following command: `scsictl -m queue_depth=21 /dev/rdisk/$dsksf`
where `/dev/rdisk/$dsksf` is the device node.
4. To make a global change to the queue depth, use the HP System Administration Manager (SAM) to edit the kernel parameter so that it equals `scsi_max_qdepth`.

Configuring the storage unit for clustering on the HP-UX 11iv2 operating system

This section describes how to configure a storage unit for clustering on the HP-UX 11iv2 operating systems that use the MC/ServiceGuard 11.14.

The steps to configure MC/ServiceGuard with the storage unit are the same as the steps in the Hewlett-Packard high availability documentation. You can find that documentation at www.docs.hp.com/hpux/ha/index.html.

After you configure your host for normal operating system access, the storage unit acts as a normal disk device in the MC/ServiceGuard configuration. Create volume groups that contain the volumes by using the Hewlett-Packard logical volume manager. This method of disk management is more reliable, easier, and more flexible than whole-disk management techniques.

When you create volume groups, you can implement PV-Links, Hewlett-Packard's built-in multipathing software for highly available disks such as the storage unit.

1. Create the volume group, using the path to the volumes that you want as the primary path to the data.
2. Extend the volume group with the path to the volumes that are intended as alternate paths.

The logical volume manager reads the label on the disk and knows that it is an alternate path to one of the volumes in the group. The logical volume manager labels the volume.

As an example, assume that you have a host that has access to a volume on a storage unit with the device nodes `c2t0d0` and `c3t0d0`. You can use the `c2` path as the primary path and create the volume group that uses only the `c2t0d0` path.

3. Extend the volume group to include the `c3t0d0` path. When you issue a **`vgdisplay -v`** command on the volume group, the command lists `c3t0d0` as an alternate link to the data.

Chapter 7. Attaching to an IBM iSeries host

This topic describes the host system requirements to attach a storage unit to an IBM iSeries host system. This topic also describes the procedures for attaching to an IBM iSeries host system.

Attaching with fibre-channel adapters to the IBM iSeries host system

This section describes the host system requirements and provides the procedure to attach a storage unit to your IBM iSeries host system with fibre-channel adapters.

Supported fibre-channel adapter cards for IBM iSeries hosts

This section lists the fibre-channel adapter cards that are supported when you attach a storage unit to an IBM iSeries host system.

- Feature code 2766
- Feature code 2787

Fibre-channel attachment requirements for IBM iSeries hosts

This topic provides fibre-channel attachment requirements for IBM iSeries hosts.

Use the following requirements when attaching the storage unit to your host system:

1. You can obtain documentation for the IBM iSeries host system from publib.boulder.ibm.com/pubs/html/as400/infocenter.htm.
2. Check the LUN limitations for your host system. See Table 6 on page 70.
3. Either you or an IBM service support representative must install the storage unit.

Note:

1. You cannot use SDD on the IBM iSeries host system.
2. With i5/OS Version 5 Release 3, you can assign storage unit LUNs to multiple iSeries fibre-channel adapters through switches, direct connection, or through a fabric. These multiple adapters must all be within the same i5/OS LPAR.

Fibre-channel attachment considerations for IBM iSeries hosts

This section provides fibre-channel attachment considerations for IBM iSeries hosts.

Note: For a list of open-systems hosts, operating systems, adapters, switches, and fabric connections that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

The storage unit creates LUN serial numbers that are eight characters in the format 0LLLLNNN, where:

LLLL A unique volume number that the storage unit assigns when it creates the LUN.

NNN The low-order, three characters of the storage unit serial number or unique three-character value that are entered using the menu option on a storage unit service panel

Note:

1. You can specify 1 - 32 LUNs for each attachment to an IBM iSeries fibre-channel adapter.
2. Fibre-channel attached LUNs are identified as the storage unit device type of 2107, on the IBM iSeries host system.
3. You can place the IBM iSeries volumes in the storage unit storage arrays according to the selected host system attachment type.
4. With a pre-i5/OS Version 5 Release 3 iSeries system, you cannot share an IBM iSeries volume with more than one fibre-channel system attachment.
5. With an i5/OS Version 5 Release 3 iSeries system, you can share an IBM iSeries volume with multiple fibre-channel system attachment. All fibre-channel adapters must be within the same i5/OS LPAR.
6. The attachment type and the available storage unit storage array capacity determine the number of volumes that you can create.
You can create 1 - 32 LUNs for a fibre-channel attachment.

Figure 37 shows an example of the display for the hardware service manager (HSM) auxiliary storage hardware resource detail for the 2766 adapter card. This same information is displayed for the 2787 adapter card.

```

Description.....: Multiple Function IOA
Type-Model.....: 2766-001
Status.....: Operational
Serial number.....: 10-22036
Part number.....: 0000003N2454
Resource name.....: DC18
Port worldwide name.....: 10000000C922D223
PCI bus.....:
  System bus.....: 35
  System board.....: 0
  System card.....: 32
Storage.....:
  I/O adapter.....: 6
  I/O bus.....:
  Controller.....:
  Device.....:

```

Figure 37. Example of the display for the auxiliary storage hardware resource detail for a 2766 or 2787 adapter card

Host limitations for IBM iSeries hosts

This topic list some host limitations for IBM iSeries hosts

See Table 6 for a description of the LUN assignments for the IBM iSeries host system.

Table 6. Host system limitations for the IBM iSeries host system

Host system	LUN limitation assignments per target
IBM iSeries (fibre-channel). You can attach the IBM iSeries through fibre-channel adapter feature 2766.	0 - 32
IBM iSeries (fibre-channel). You can attach the IBM iSeries through fibre-channel adapter feature 2787.	0 - 32

IBM iSeries hardware

This section describes the hardware that you can use with IBM iSeries hosts.

The DS8000 supports the following models for the IBM iSeries hosts:

- Models 270, 800, 810, 820, 825, 830, 840, 870, 890
- The IBM eServer i5 product line consists of model 550, model 520, and model 570.

For more information on supported switches, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>

The IBM iSeries fibre-channel adapter automatically detects the attachment protocol. You do not need to manually perform a system configuration.

The iSeries fibre-channel adapter operates at 2 Gb on OS/400 Version 5 Release 2 or i5/OS Version 5 Release 3.

IBM iSeries software

This section describes the software that you can use with IBM iSeries hosts.

Before you attach a storage unit to the IBM iSeries host using a fibre-channel adapter, you must install one of the following operating systems:

- i5/OS Version 5 Release 3
- OS/400 Version 5 Release 2
- Red Hat Enterprise Linux 3.0
- SUSE SLES 9

For more information on supported operating systems, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>

General information for configurations for IBM iSeries hosts

This section contains general configuration information for IBM iSeries hosts.

The following list identifies some general information about connecting the IBM iSeries host through a switch:

- The IBM iSeries supports only a homogeneous environment (only IBM iSeries initiators). You can establish a homogeneous environment by using the logical zoning of the switch. All host systems within an IBM iSeries zone must be IBM iSeries systems.
- CNT: A list of supported environments (servers, operating systems, adapters) and hardware and software prerequisites for the CNT FC/9000 is available at <http://www.storage.ibm.com/ibmsan/products/directors/index.html>.
- Brocade: A list of supported environments (servers, operating systems, and adapters) and hardware and software prerequisites for the Brocade 2109 M12 and 2109 M14 Fibre Channel Director is available at <http://www.brocade.com>.
- McDATA: A list of supported environments (servers, operating systems and adapters) and hardware and software prerequisites for the McDATA ED-6064 is available at <http://www.storage.ibm.com/ibmsan/products/2032/6064/>.
- Cisco: A list of supported environments (servers, operating systems, and adapters) and hardware and software prerequisites for the Cisco is available at <http://www.cisco.com/ibmsan/cisco/index.html>.

Recommended configurations for IBM iSeries hosts

This topic provides recommended configurations for IBM iSeries hosts.

You can use the following configurations for each feature code:

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

Figure 38 shows an example of the display for the HSM logical hardware resources associated with the IOP.

Opt	Description	Type-Model	Status	Resource Name
	Combined Function IOP	2843-001	Operational	CMB04
	Storage IOA	2766-001	Operational	DC18
	Disk Unit	2107-A82	Operational	DD143
	Disk Unit	2107-A81	Operational	DD140
	Disk Unit	2107-A81	Operational	DD101

Figure 38. Example of the logical hardware resources associated with an IOP

Figure 39 shows an example of the display for the HSM auxiliary storage hardware resource detail for the storage unit.

Description.....	Disk unit
Type-Model.....	2107-A82
Status.....	Operational
Serial number.....	75-1409194
Part number.....	
Resource name.....	DD143
licensed machine code.....	FFFFFFFF
Level.....	
PCI bus.....	
System bus.....	35
System board.....	0
System card.....	32
Storage.....	
I/O adapter.....	6
I/O bus.....	0
Controller.....	1
Device.....	1

Figure 39. Example of the display for the auxiliary storage hardware resource detail for the storage unit

You can define the storage unit LUNs as either protected or unprotected. From a storage unit physical configuration view point, all IBM iSeries volumes are RAID-5 or RAID-10 volumes and are protected within the storage unit. When you create the IBM iSeries LUNs by using the IBM TotalStorage DS Storage Manager, you can create them as logically protected or unprotected. Table 7 shows the disk capacity for the protected and unprotected models. To logically unprotect a storage LUN, allow the iSeries host to perform remote load source mirroring to that device. Because the load source is mirrored on an external LUN, the storage unit can copy or transfer this load source as a disaster recovery backup. When you use the iSeries tools kit, an iSeries host in a remote location, using a copy of the original load source, can recover this load source and start running as if this recovery box was the original source host.

Table 7. Capacity and models of disk volumes for IBM iSeries

Size	Type	Protected model	Unprotected model	Release support
8.5 GB	2107	A01	A81	Version 5 Release 2 or later
17.5 GB	2107	A02	A82	Version 5 Release 2 or later
35.1 GB	2107	A05	A85	Version 5 Release 2 or later
70.5 GB	2107	A04	A84	Version 5 Release 2 or later
141.1 GB	2107	A06	A86	Version 5 Release 3 or later
282.2 GB	2107	A07	A87	Version 5 Release 3 or later

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

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

To improve the availability of fibre-channel attached disks (when there is more than one 2766 or 2787 I/O adapters), place each I/O adapter and its I/O processor as far as possible from the other fibre-channel disk I/O adapters or I/O processors in a system. If possible, place them on different HSL loops, different towers, or on different multiadapter bridges.

With i5/OS Version 5 Release 3, path information is available for the Disk Unit hardware configuration from iSeries Navigator.

Running the Linux operating system on an IBM i5 server

The Linux operating system can run on an IBM i5 server.

Disabling automatic system updates

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

If your host is attached to the DS8000 series and is running the IBM Multipath Subsystem Device Driver (SDD), you should consider turning off this automatic update feature because some drivers supplied by IBM, like SDD, are dependent on a specific kernel and cannot function in the presence of a new kernel. Similarly, host bus adapters drivers must be compiled against specific kernels to function optimally. By allowing automatic update of the kernel, you risk an unexpected impact to your host system.

Supported fibre-channel adapters for IBM i5 servers running the Linux operating system

Fibre-channel adapters are supported for IBM i5 servers running the Linux operating system.

The following adapter cards are supported for the IBM i5 server when it is running the Linux operating system:

- Feature code 0612-designates a 2766 assigned to Linux
- Feature code 0626-designates a 5704 assigned to Linux
- Feature code 0646-used on IBM i5 servers

Running the Linux operating system in a guest partition on an IBM i5 servers

This topic provides instructions for running the Linux operating system on IBM i5 servers.

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

- Plan to run Linux as a hosted or nonhosted guest partition

Find out what you need to do before you install Linux on the i5 server.

Understand what software and hardware requirements that you need to support Linux. Find out the configuration options that are available and which options fit your company's needs. Find out if the system that you own requires you to disable the multitasking function of your server's processor.

- Create a guest partition to run Linux

Understand how to configure a guest partition using system service tools (SST) and how to install Linux on the i5 server. You can also find information about supported I/O adapters (IOAs) and how to configure a network server description (NWSD.)

- Manage Linux in a guest partition

Find the information you need to manage a guest partition running Linux. Use the information to understand what IPL types you can use and how Linux partitions can communicate with other partitions on the server.

- Ordering a new server or upgrading an existing server

Use the LPAR Validation tool to validate that your planned partitions are valid. You can also find the contact information that you need to order a new server.

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

- <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>
- <http://publib.boulder.ibm.com/series/v5r2/ic2924/index.htm>

Planning to run Linux in a hosted or nonhosted guest partition

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

Find out what you need to do before you install Linux on the i5 server. Understand what software and hardware requirements are needed to support Linux. Find out the configuration options that are available and which options fit your company's needs.

Linux support on iSeries servers

Evaluate each i5 server to determine if your hardware supports Linux. To successfully partition an i5 server to run Linux, the server requires specific hardware and software. The primary partition must run i5/OS Version 5 Release 3 and be updated with the latest programming temporary fixes. You can find the latest Linux related i5/OS programming temporary fixes at <http://publib.boulder.ibm.com/series/v5r2/ic2924/index.htm>.

Linux is not supported on a primary partition.

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

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

To determine whether or not your system will support Linux in a guest partition, go to <http://publib.boulder.ibm.com/series/v5r2/ic2924/index.htm>.

Creating a guest partition to run Linux

This sections helps you understand how to configure a guest partition and how to install Linux on the i5 server.

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

Hosted versus nonhosted guest partition running Linux

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

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

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

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

Obtaining Linux for i5 servers

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

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

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

iSeries I/O adapters (IOAs) supported by Linux

You can assign IOAs to a guest partition. See publib.boulder.ibm.com/series/v5r2/ic2924/index.htm for a list of adapters that the i5

server supports in a guest partition running Linux.

Managing Linux in a guest partition

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

Virtual I/O in a guest partition running Linux

Virtual I/O resources are devices that are owned by the hosting i5/OS partition. The i5 Linux kernel and i5/OS support several different kinds of virtual I/O resources. They are virtual console, virtual disk unit, virtual CD, virtual tape, and virtual Ethernet.

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

When you use virtual disk for a Linux partition, the i5/OS partitions control connectivity to the real disk storage. The i5/OS hosting partition and its operating system versions control the storage unit connectivity solely in this configuration. For more information about i5/OS connectivity to the storage unit, see <http://www-1.ibm.com/servers/storage/disk/ds8000>.

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

For more detailed information about how your company might use a guest partition with I/O resources go to <http://publib.boulder.ibm.com/series/v5r2/ic2924/index.htm>.

Directly attached I/O in a guest partition running Linux

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

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

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

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

For more detailed information about how your company might use a guest partition with I/O resources go to <http://publib.boulder.ibm.com/series/v5r2/ic2924/index.htm> or contact your IBM marketing representative or IBM Business Partner for further assistance on using directly attached I/Os in a guest partition.

Ordering a new server or upgrading an existing server to run a guest partition

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

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

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

Chapter 8. Attaching to an IBM NAS Gateway 500 host

This chapter provides instructions, requirements, and considerations for attaching an IBM NAS Gateway 500 host to a storage unit.

Supported adapter cards for IBM NAS Gateway 500 hosts

This topic lists the supported host adapters for the IBM NAS Gateway 500 host.

The following fibre-channel adapter cards are supported for the IBM NAS Gateway 500:

- Feature code 6239
- Feature code 6240

Feature code 6239 is a 1-port 2 gigabit fibre-channel host adapter card. Feature code 6240 is a 2-port 2 gigabit fibre-channel host adapter card.

Finding the worldwide port name

You can obtain the worldwide port names (WWPN) of the fibre-channel adapters installed on NAS Gateway 500 through a Web browser or through the DS CLI. The following sections provide instructions for using these methods.

Obtaining WWPNs using a Web browser

This section provides instructions for obtaining WWPNs using a web browser.

If your external storage requires you to enter a worldwide port name (WWPN) for the fibre-channel host adapters that are installed in your NAS Gateway 500, you can obtain these WWPNs by using the Internet.

1. Open a Web browser.
2. Enter the following Web address: `http://hostname/NAS500GetWWN.html` where *hostname* is the host name or IP address of your NAS Gateway 500 system. If your NAS Gateway 500 is not within the same IP subnet, use the fully qualified domain name that you use with the DNS name resolution. For example: `nasgateway500.servers.mycompany.com`.

Obtaining WWPNs through the command-line interface

This section provides instruction for obtaining WWPNs through the command line interface.

1. Login to the NAS Gateway 500 using the root user ID from a serial terminal.
2. Run the following command to install the WWPNs of all fibre-channel adapters:

```
lscfg -vp1 "fcs*" |grep Network
```

Figure 40 on page 80 shows an example of the output that you would receive.

```
(/)-->lscfg -vpl "fcs*" |grep Network
Network Address.....10000000C93487CA
Network Address.....10000000C934863F
Network Address.....10000000C93487B8
Network Address.....10000000C934864F
```

Figure 40. Example output from the `lscfg -vpl "fcs*" |grep Network` command.

3. You can optionally use the following command to put all vital product data of the fibre-channel adapters installed on NAS Gateway 500 into a text file.

```
lscfg -vpl "fcs*" > foo.txt
```

You can keep this file for future use. Figure 41 shows an example of the text output that you would receive. You can find information for your WWPN, location information, microcode level, part number, other information about your fibre-channel adapters in this file.

```
fcs2 U0.1-P2-I6/Q1 FC Adapter
Part Number.....00P4295
EC Level.....A
Serial Number.....1E323088E2
Manufacturer.....001E
Feature Code/Marketing ID...5704
FRU Number..... 00P4297
Device Specific.(ZM).....3
Network Address.....10000000C93487CA
ROS Level and ID.....02E01035
Device Specific.(Z0).....2003806D
Device Specific.(Z1).....00000000
Device Specific.(Z2).....00000000
Device Specific.(Z3).....03000909
Device Specific.(Z4).....FF601032
Device Specific.(Z5).....02E01035
Device Specific.(Z6).....06631035
Device Specific.(Z7).....07631035
Device Specific.(Z8).....20000000C93487CA
Device Specific.(Z9).....HS1.00X5
Device Specific.(ZA).....H1D1.00X5
Device Specific.(ZB).....H2D1.00X5
Device Specific.(YL).....U0.1-P2-I6/Q1
fcs3 U0.1-P2-I5/Q1 FC Adapter
Part Number.....00P4295
EC Level.....A
Serial Number.....1E3230890F
Manufacturer.....001E
Feature Code/Marketing ID...5704
FRU Number..... 00P4297
Device Specific.(ZM).....3
Network Address.....10000000C934863F
ROS Level and ID.....02E01035
Device Specific.(Z0).....2003806D
Device Specific.(Z1).....00000000
Device Specific.(Z2).....00000000
```

Figure 41. Example output saved to a text file

Multipathing support for NAS Gateway 500

Depending on the storage subsystem that you have, you can choose the multipathing software. Be sure that your configuration needs multipathing, because the configuration might not be supported to run multipathing software with only one path.

Multipath I/O and SDD considerations for NAS Gateway 500

This is an overview of multipath I/O and SDD considerations for NAS Gateway 500.

The NAS Gateway 500 is shipped with the following preinstalled multipathing drivers:

- Multipath I/O (MPIO)
- Multipathing Software for FASiT storage Subsystems (RDAC)

MPIO is a technology that was introduced as part of the AIX 5.2 operating system. If you want to connect to a storage unit using multiple paths, you must use Subsystem Device Driver (SDD). The MPIO driver also performs this connection for an unclustered NAS Gateway 500 attached to storage unit. SDD (multipathing software for the storage unit) is not preinstalled on the NAS Gateway 500. However, you can connect FASiT storage to your NAS Gateway 500 without adding additional software.

You can download the latest version of SDD from the Internet. To check for the latest versions of SDD, go to the following website: <http://www-1.ibm.com/servers/storage/support/software/sdd.html>. You must verify that the provided SDD version supports the NAS Gateway 500 and must check the SDD prerequisites in the readme files and in the installation documentation. Restrictions and prerequisites can appear. Look for any required program temporary fixes (PTFs). You can also verify prerequisites regarding the DS8000 Interoperability Matrix at the following web site: <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>. Look for entries regarding the NAS Gateway 500.

Host attachment multipathing scripts

This section provides instructions for applying host attachment multipathing scripts.

Before installing the multipathing software SDD, you must install storage unit attachment scripts on the NAS nodes.

1. Prepare your environment for migration from a MPIO environment to a SDD environment.
 - a. If you install the host attachment scripts (**devices.fcp.disk.ibm.rte**) for the first time on the NAS Gateway 500, you must remove the MPIO drivers. To check if the attachment scripts are installed, enter the following command on a command line: `ls1pp -La | grep 2107`
 - b. Figure 42 on page 82 displays the situation where MPIO (**devices.fcp.disk.ibm.mpio.rte**) is installed, and the 2107 host attachment scripts (**devices.fcp.disk.ibm.rte**, a prerequisite for SDD) is not installed. Both versions will not coexist on the NAS Node. If they both exist, you must unconfigure and remove the MPIO package.

```

</>--->lslpp -La | grep ibm2107
      devices.fcp.disk.ibm.mpio.rte      1.0.0.0   C   F   IBM 2107 Host Attachment
</>--->_

```

Figure 42. Example of installed file set

2. Perform these steps to unconfigure MPIO devices and decide which disks to remove:
 - a. Issue the following command, as a root user, from the command prompt:
`lsdev -Cc disk.`
Attention: In the following example, the first two disks (16 Bit LVD SCSI Disk Drive) **hdisk0** and **hdisk1** are the operating system disks (internal SCSI disks) You must not remove them. Be aware that these numbers are only examples and can be different on your system.
 - b. Use the following command to delete all MPIO disks: `rmdev -d1 hdiskX -R`. This command deletes the disk from the ODM including all child devices. For example, `rmdev -d1 hdisk2 -R`
 Figure 43 shows an environment with only the internal disks. Be aware that the disk numbers on your system can be different.

```

</>--->lsdev -Cc disk
hdisk0 Available 1S-08-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 1S-08-00-9,0 16 Bit LVD SCSI Disk Drive

```

Figure 43. Example environment where all MPIO devices removed

- c. Now you can remove the **devices.fcp.disk.ibm.mpio.rte** file set. Open a command line, as root user, and enter the following command: `installp -u devices.fcp.disk.ibm.mpio.rte`. This command removes the MPIO file set.
 The NAS node is now prepared for the **devices.fcp.disk.ibm.rte** installation.
3. Find the appropriate scripts on a CD that comes with your storage unit. You can ask your storage unit administrator to get the CD or file sets. To download the file sets from the Internet, go to the following Website:
<http://www-1.ibm.com/servers/storage/support/software/sdd.html>
4. Follow the link named **2107 host attachment scripts**, then select **2107 Host Attachment Script for SDD on AIX**. This selection takes you to a Web page where you can select **AIX Host Attachment Script** under the Download column.
5. Download the **devices.fcp.disk.ibm.rte.tar** file from the Web page into the following local directory on your NAS node: **/usr/sys/inst.images**
6. Change to this directory by using the `cd /usr/sys/inst.images` command. Extract the downloaded `devices.fcp.disk.ibm.rte.tar` file with the `tar -xvf devices.fcp.disk.ibm.rte.tar` command.
7. Create a new table of content file (.toc) by entering the following command:
`inutoc .` (this period is part of the command). This file is used later by the installation procedure.
8. Run the install SMIT menu (with the fastpath) by entering the following command: **smitty install_latest**. Specify the path (you can use a period (.))

because we changed the directory before, or you can specify the full path). Press **Esc+4** and select the appropriate input directory. For example, **/usr/sys/inst.images <Installation Directory>**.

9. Place the cursor at **SOFTWARE** and press **Esc+4** to list the software to install.
10. Move the cursor to **ibm2107** and press **Esc+7** to select the IBM 2107 disk device (the > should appear in front of the **ibm2107** software).
11. Press Enter and the installation begins.
12. Use the following command to verify that the file set is installed: `lslpp -La | grep ibm2107`, as shown in Figure 44.

```
</usr/sys/inst.images>-->lslpp -La | grep ibm2107
devices.fcp.disk.ibm.rte      32.6.100.18   C    F    IBM 2107 Disk Device
</usr/sys/inst.images>-->_
```

Figure 44. Example of installed host attachment scripts

If the storage unit and SAN is correctly set up and volumes are assigned to the NAS Gateway 500, you might see a listing of disks when you issue the `lsdev -Cc disk` command. All disks residing on the storage unit should display the following information in the description: `hdiskx..... FC2105xxx`.

The system disks `hdisk0` (and `hdisk1`, if the system has two internal disks) do not have the 2107 description.

Tip: At this point, it is not advised to configure volume groups, logical volumes, or file systems on storage unit disks. This should be done after SDD is installed. If volume groups are created before, the administrator has to convert the `hdisk` volume group to a `vpath` volume group. (see the `hd2vp` command in the SDD documentation).

If the physical volumes on an SDD volume group physical volumes are mixed with `hdisk` devices and `vpath` devices, you must run the `dpovgfix` utility to fix this problem. Otherwise, SDD does not function properly. Issue the `dpovgfix vg_name` command to fix this problem.

Multipathing with the Subsystem Device Driver

This section provides instructions for installing the multipathing software for SDD.

Ensure that you have followed the steps for preparing and installing the host attachment scripts. See “Host attachment multipathing scripts” on page 81 for specific instructions.

SDD drivers are provided with the storage unit, but it is advisable to look for the latest version on the Internet. The SDD file sets (nonconcurrent HACMP), readme files and documentation can be found at the following Web site:

<http://www-1.ibm.com/servers/storage/support/software/sdd.html>. More information about installing and using SDD, see the *IBM TotalStorage Multipath Subsystem Device Driver User's Guide*.

1. To verify that you have the latest version of SDD, look on the Internet for the latest version.
 - a. From the Subsystem Device Driver web page, go to the Download section and select **Subsystem device drivers**.
 - b. Choose the appropriate version. Remember that NAS Gateway is AIX powered, so choose **Subsystem device driver for AIX**.

- c. Check the documentation, the “readme” file, and the Web site for prerequisites concerning SDD.
- d. You can download the file set to the NAS node with the file transfer program (FTP). Then extract the file set as shown in Figure 45.

```
</usr/sys/inst.images>-->tar -xv -f devices.sdd.52.rte.tar
x devices.sdd.52.rte, 1177600 bytes, 2300 media blocks.
</usr/sys/inst.images>-->_
```

Figure 45. Example of extracting the SDD archive

- e. Change the directory to `cd /usr/sys/inst.images` and run the `inutoc . exec` to create a new **.toc** file in this directory.
If you are going to upgrade SDD, refer to the *Subsystem Device Driver User's Guide the IBM Enterprise Storage Server and the IBM SAN Volume Controller* for details on upgrading to SDD on AIX 5.2.
2. The following installation guidelines are for a fresh SDD installation only:
 - a. Ensure that you are logged in as a root user.
 - b. Issue the following command at a command prompt: `smitty install_latest`.
 - c. Choose a period (.) if you changed the directory to `/usr/sys/inst.images` before. Otherwise, select the directory and press **Esc+4** to get to the SMIT install menu.
 - d. Place the cursor at **SOFTWARE** to install and press **ESC+4**.
 - e. Use the up and down keys to place the cursor to **devices.sdd.52** and press **ESC+7** to select this SDD file set. The **>** sign shows the selected item.
 - f. Press Enter again to return to the SMIT installation menu.

Note: If you want to do a preview first, select yes in the **PREVIEW only?** field.

- g. Proceed with the installation by pressing Enter. SMIT will ask you to confirm the installation by pressing Enter again. The installation should finish successfully.
- h. Verify that the file set is installed by issuing the `ls1pp -La | grep sdd` command. Ensure that the storage unit is correctly configured and attached, and that zoning in your SAN environment is correct. Now you can run the configuration Manager (**cfgmgr**) on a NAS Gateway 500 command line.

Attention: Only use the **cfgmgr** command that is located in the `/opt/nas/bin/cfgmgr` directory.

If everything is correctly installed, you will see 2107 disks (hdisks and vpath devices with FC2105xxx descriptions) when you run the `lsdev -Cc disk` command. You can also try the following SDD commands: `datapath query adapter`, or `datapath query device`, or `lsvpcfg`.

After SDD is successfully installed, you can configure VGs, LVs and file systems. Keep in mind that SDD provides new commands like `mkvg4vp` and `extenvg4vp`.

Note: You should have only PVIDs on vpath devices (`lsdev -Cc disk`).

Chapter 9. Attaching to an IBM RS/6000 or IBM eServer pSeries host

This chapter provides instructions, requirements, and considerations for attaching an IBM RS/6000 or IBM eServer pSeries host to a storage unit.

Installing the 2107 host attachment package on IBM pSeries AIX hosts

This section provides the instructions to install the host attachment package for the storage unit. Run the host attachment package on each AIX host system that is attached to the storage unit.

Preparing for installation of the host attachment package on IBM pSeries AIX hosts

This section describes preparing for installation of the host attachment package on IBM pSeries AIX hosts.

1. Attach the storage unit to your host system. See “General information about attaching to open-systems host with fibre-channel adapters” on page 16.
2. Turn on the host system and all attachments.
3. Ensure that you have root access.
4. Ensure that you have administrator knowledge.
5. Ensure that you have knowledge of the System Management Interface Tool (SMIT).

Note: You are not required to vary off the volume group before you install the new version.

Installing the host attachment package on IBM pSeries AIX hosts

This section provides instructions for installing the host attachment package on IBM pSeries AIX hosts.

You must have superuser authority to complete the instructions.

Note: The following procedure is an example. The example uses `/dev/cd0` for the address of the CD. Your address might be different.

1. From your desktop window, type `smit install_update` to go directly to the installation panel.
2. Click **Install Software** and press Enter.
3. Press F4 to open the Input Device/Directory for Software window.
4. Select the CD drive that you are using for the installation, for example, `/dev/cd0`.
5. Press Enter. The Install and Update from the Latest Available Software window opens.
6. Click **Software to Install** and press F4.
7. Select **Software Packages** and press F7.
The software package name is `devices.fcp.disk.ibm.rte`.
8. Press Enter. The Install and Update from the Latest Available Software panel opens again with the name of the selected software highlighted.
9. Check the default option settings to ensure that they are what you need.

10. Press Enter to install the software. SMIT responds with the following question:
Are you sure?
11. Press Enter to continue. The installation process might take several minutes. A message displays when the installation process is complete.
12. Press F10 when the installation process is complete.
13. Exit from SMIT.
14. Remove the compact disc.
15. Restart the host system or unconfigure and reconfigure the disks. You must restart the host system or unconfigure and reconfigure the disks before the changes can take effect.

Upgrading the host attachment package on IBM pSeries AIX hosts

You can upgrade the host attachment package on an IBM pSeries hosts.

For current information about upgrading the host attachment package on an IBM pSeries AIX hosts, see the *IBM TotalStorage Multipath Subsystem Device Driver User's Guide*.

Attaching with fibre-channel adapters

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

Supported fibre-channel adapter cards for IBM pSeries hosts

The following fibre-channel adapters are supported for RS/6000 or pSeries host systems:

- Feature code 6228
- Feature code 6239
- Feature code 5716

This section also describes the host system requirements and provides the procedures to attach a storage unit to the following host systems:

- RS/6000
- pSeries
- eServer p5 with fibre-channel adapters
- RS/6000 ScalablePOWER Parallel (SP) Complex
- pSeries SP Complex

Note: For an RS/6000 or pSeries host system, you can use either of the following topologies:

- Point-to-point (switched fabric) topology
- Arbitrated loop topology

The RS/6000 and pSeries host systems *do not support* more than one host bus adapter on the loop. The RS/6000 and pSeries host systems support a direct connection of the RS/6000 and pSeries host systems to a storage unit by using the fibre-channel arbitrated loop topology.

For information about running MPIO on a pSeries host system, or for a list of open-systems hosts, operating systems, adapters and switches that IBM supports,

see the *Interoperability Matrix* at
<http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Fibre-channel attachment requirements for IBM pSeries hosts

This section list the fibre-channel attachment requirements for IBM pSeries hosts.

This section lists the requirements for attaching the storage unit to your host system:

- Ensure that you have the installation script files. These files are on the CD that you receive with the storage unit.
- Ensure that you have 1 MB minimum of hard disk space available to install the AIX host attachment package.
- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *User's Guide* is on the TotalStorage CD that you receive with the storage unit.
- For details about the release level for your operating system, see the *Interoperability Matrix* at
<http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Either you or an IBM service support representative (SSR) must perform the following tasks to install and configure a storage unit:

1. The IBM SSR installs the storage unit by using the procedures in the *IBM TotalStorage DS8000 Service Guide*.
2. Either you or an IBM SSR defines the fibre-port configuration if you did perform the installation of the storage unit or fibre-channel adapters.
3. Either you or an IBM SSR configures the host system for the storage unit. Use the instructions in your host system publications.
4. Either you or an IBM SSR checks the LUN limitations for the RS/6000 and pSeries.

Note: The SDD supports RS/6000 and pSeries host systems in a clustering environment. To have failover protection on an open system, SDD requires at least two paths, and SDD allows a maximum of 32 paths.

Fibre-channel attachment considerations for IBM pSeries hosts

This section contain information about fibre-channel attachment considerations for IBM pSeries hosts.

For up to date information about the servers, operating systems, host adapters, and fabric connections that IBM supports, see the *Interoperability Matrix* at
<http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

The storage unit supports eServer BladeCenter JS20 with blade 8842 when operating in an IBM pSeries, 32-bit environment.

Verifying the configuration of the storage unit for fibre-channel adapters on the AIX host system

This topic provides information to verify the configuration of the storage unit for fibre-channel adapters on the AIX host system.

To verify the configuration of the storage unit on the AIX host system, type the following command: `lsdev -Cc disk | grep 2107`

Figure 46 displays example results of the command if the configuration is successful.

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

Figure 46. Example of a list of devices displayed when you use the `lsdev -Cc disk | grep 2107` command for fibre-channel

Figure 47 displays example results of the `lsdisk` command to list all the disks in the systems which do not contain 2107 disk devices.

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

Figure 47. Example of a list of other devices displayed when you use the `lsdisk` command for fibre-channel.

Making SAN changes for IBM pSeries hosts

This topic describes SAN changes that you need to make for an IBM pSeries hosts.

Before you change the switch port connected to the storage unit or the switch domain ID, perform the steps below to remove the fibre-channel adapter:

The 5716, 6228, and 6239 host adapter cards do not support dynamic SAN configuration changes. Do not change the domain ID of the switch port of the switch connected to the storage unit; otherwise, you might lose the connection to the storage unit.

1. Run the **umount** command to unmount the files on the file system.
For example, type `umount x`, where *x* is the file system name. If you are on an SP host, go to step 2, otherwise go to step 4.
2. For SP environments, type `suspendvsd` to suspend the virtual shared disks with the volume groups.
3. For SP environments, type `stopvsd` to stop the virtual shared disks that are associated with the volume groups.
4. Run the **varyoffvg** command for the 2107 volume group.
For example, type `varyoffvg VGname`
5. Type `rmdev -dl fcs(x) -R` on the command line to remove the fibre-channel adapter. The value of *x* represents the fibre-channel number that you want to remove.
6. Type `cfgmgr` to reinstall the adapter and rediscover the 2107 hdisks.
7. Type `varyonvg` to vary on the volume groups.
8. Remount the file system.

Support for fibre-channel boot

The following sections describe the prerequisites, the types of IBM RS/6000 or IBM pSeries hosts that are supported, and required level of adapter microcode.

Prerequisites for setting up the IBM pSeries host as a fibre-channel boot device

This section contain prerequisites for setting up the IBM pSeries host as a fibre-channel boot device

Note: RS/6000 supports fibre-channel boot with or without SDD. However, if you are using SDD with RS/6000 SP, you must install APAR IY37451.

Before you can set up the RS/6000 or IBM pSeries as a boot device, you must meet the following prerequisites:

- An IBM pSeries or IBM RS/6000 system with the correct level of firmware
For the most current list of IBM RS/6000 or IBM pSeries hosts that are supported, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
- Feature code 5716 at the correct microcode level and requires AIX 5.2H, AIX 5.3, or higher.
- Feature code 6228 at the correct microcode level.
- Feature code 6239 needs to be at the correct microcode level and requires AIX 5.1 ML04 and above.
- AIX 5.1
For the SP, PSSP must be at 3.2 with APAR IY22732 (included in PSSP 3.2 PTF 14).
- A copy of the **devices.fcp.disk.ibm.rte** file.

For the most current microcode and firmware levels, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Fibre-channel boot considerations for IBM pSeries hosts

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

Before you plan and use fibre-channel, consider the following items:

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

Supported IBM RS/6000 or IBM pSeries hosts for fibre-channel boot

This section lists supported IBM RS/6000 or IBM pSeries hosts for fibre-channel boot.

For up to date information about the hosts for fibre-channel boot that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Supported levels of firmware for fibre-channel boot on IBM pSeries hosts

This topic provides information about the supported levels of firmware for fibre-channel boot on IBM pSeries hosts.

For information about the most up-to-date levels of firmware, go to techsupport.services.ibm.com/server/mdownload/download.html.

Supported levels of fibre-channel adapter microcode on IBM pSeries hosts

This topic provides supported levels of fibre-channel adapter microcode on IBM pSeries hosts.

For more information about the required level of microcode for your adapter, go to <http://techsupport.services.ibm.com/server/mdownload>

Note: You must get the microcode from an IBM Web site.

Installation mechanisms that PSSP supports for boot install from fibre-channel SAN DASD on IBM pSeries hosts

This topic provides installation mechanisms for boot install from fibre-channel SAN DASD on IBM pSeries hosts.

You can direct a PSSP to do a boot install on an IBM RS/6000 SP node from the following items:

- SAN_DISKID which is derived from the worldwide port name and LUN ID
This is the preferred mechanism.
- PVID
- HDISK

Support for disk configurations for RS/6000 for a fibre-channel boot install

This section lists the methods that you can use to do a disk configuration for the IBM RS/6000 for a boot install from fibre-channel Storage Area Networks (SAN) DASD.

- Mirrored volume
- Alternate volume
- Individual volume

Support for fibre-channel boot when a disk subsystem is attached on IBM pSeries hosts

This section provides support information for fibre-channel boot when a disk subsystem is attached on IBM pSeries hosts.

For a list of open systems hosts, operating systems, adapters and switches that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

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

This section provides the instructions to attach one or two storage units to multiple host systems without the High Availability Cluster Multi-Processing (HACMP™) clustering solution.

Considerations for attaching to multiple RS/6000 or pSeries hosts without the HACMP host system

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

Install HACMP to define and access a unique journaled file system (JFS) file stored on a single storage unit from any attached host system.

When attaching multiple host systems to a storage unit, consider the following:

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

Saving data on the storage unit when attaching multiple RS/6000 or pSeries host systems to the storage unit

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

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

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

Restoring data on the storage unit when attaching multiple RS/6000 or pSeries host systems to the storage unit

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

This procedure assumes that you have preserved the data by following the instructions in “Saving data on the storage unit when attaching multiple RS/6000 or pSeries host systems to the storage unit” on page 91.

1. Check to be sure that the storage unit physical volumes (hdisks) are available.
Type `lsdev -Cc disk` to display the hdisks on the host system.
2. Type `importvg xxxx` where `xxxx` is the name of the volume group to restore the storage unit volume groups to the applicable host systems that are connected to the storage unit, one system at a time.
3. Type `smit chvg` to verify that **No** is selected for the **Activate Volume Group Automatically** and **A Quorum and of disks required to keep the volume group online?**
4. Type `varyonvg VGname` to vary-on the storage unit volume groups to the applicable host systems connected to the storage unit. Perform the steps one system at a time.
5. Type `mount` to mount the storage unit volume groups to the applicable host systems that are connected to the storage unit. Perform the steps one system at a time.
6. If you want to create new volume groups, go to step 3 in “Restoring data on the storage unit when attaching multiple RS/6000 or pSeries host systems to the storage unit.”

Running the Linux operating system on an IBM pSeries host

This section describes how to attach a pSeries host system running the Linux operating system with a fibre-channel adapter.

Attachment considerations for running the Linux operating system on an IBM pSeries host

A pSeries host supports the Subsystem Device Driver (SDD) when using the Linux operating system. The following sections provide information that you need to prepare your system for SDD. You can also find instructions to install and verify SDD.

Storage unit attachment is supported either by directly attaching a storage unit port to the host adapter, or by attaching through a fibre-channel switch. You must specify the proper attachment topology using the storage specialist. Select **Loop** to indicate direct attachment.

Disabling automatic system updates

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

Hosts that are attached to the DS8000 series and running the IBM Multipath Subsystem Device Driver (SDD) should consider turning off this automatic update feature. Some drivers supplied by IBM, like SDD, are dependent on a specific kernel, and will cease to function in the presence of a new kernel. Similarly, host bus adapter (HBA) drivers need to be compiled against specific kernels to function optimally. By allowing automatic updates of the kernel, you risk impacting your host systems unexpectedly.

Hardware requirements for the Linux operating system on the pSeries host

This section provides the list of hardware that you need to ensure that the Subsystem Device Driver (SDD) installs and operates successfully.

You must have the following hardware to ensure that the Subsystem Device Driver (SDD) installs and operates successfully:

- The storage unit device
- One or more pairs of fibre-channel host adapters
To use the I/O load balancing and failover features of SDD, you must have a minimum of two paths. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html> for more information about the fibre-channel adapters that you can use on your Linux host system.
- Subsystem LUNs that you created and configured for multiport access. Each LUN can have up to 32 SCSI disk instances, one for each path on the server.
- A fibre-optic cable to connect each fibre-channel adapter to a storage unit port, or to two switch ports subsequently zoned to storage unit ports.

Software requirement for the Linux operating system on the pSeries host

This section provides the list of software that you need to ensure that the Subsystem Device Driver (SDD) installs and operates successfully.

If using SUSE SLES 8 or 9 or Red Hat Enterprise Linux 3.0, you must have Service Refresh 1 to ensure that SDD installs and operates successfully.

Unsupported Environments

SDD does not support the following environments for the Linux operating system on a pSeries host:

- A host system with fibre-channel connection to a shared storage unit LUN
- Logical Volume Manager
- System start from an SDD pseudo device
- EXT3 filesystem on an SDD pseudo device
- System paging file on an SDD pseudo device
- Single-path mode during concurrent download and installation of LMC
- Single-path mode during any maintenance that impacts the path attachment, such as replacing a storage unit host adapter

Preparing to install the Subsystem Device Driver for the Linux operating system on the pSeries host

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

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

Configuring the storage unit

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

See *IBM TotalStorage DS8000 Introduction and Planning Guide* for more information about configuring the storage unit for the SDD.

Configuring the fibre-channel adapters for the storage unit

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

Installing the Subsystem Device Driver on the pSeries host running the Linux operating system

This section provides the steps you need to complete to install the SDD on your pSeries host running the Linux operating system.

Before you install the Subsystem Device Driver (SDD), make sure that you have root access to your Linux host system.

This section provides the steps you need to complete to install the SDD on the Linux operating system for your pSeries host. If you have already installed the SDD on your host system, and need to upgrade, see “Upgrading the Subsystem Device Driver for the Linux operating system on the pSeries host” on page 95.

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

Preparing for installation ...
IBMsdd-1.4.0.0-6

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

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

This section provides steps you must complete to upgrade the Subsystem Device Driver (SDD) for the Linux operating system on the pSeries host. You must have a preexisting version of the SDD installed on your host before you can perform these steps. See “Installing the Subsystem Device Driver on the pSeries host running the Linux operating system” on page 94 if you are installing the SDD for the first time.

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

Preparing for installation ...
IBMsdd-1.4.0.0-6

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

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

This section provides the steps to verify the Subsystem Device Driver (SDD) installation on the Linux operating system for your pSeries host. This procedure is applicable for either a new installation or an upgrade of the SDD.

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

Filename	Location	Description
sdd-mod.o-xxx	/opt/IBMsdd	The SDD device driver file, where xxx represents the kernel level of your host system and y represents smp or up.
vpath.conf	/etc	The SDD configuration file.
executables	/opt/IBMsdd/bin	The SDD configuration and status tools.
	/usr/bin	The symbolic links to the SDD utilities.

Filename	Location	Description
sdd.rcscript	/etc/init.d/sdd	The symbolic link for the SDD system startup option.
	/usr/sbin/sdd	The symbolic link for the SDD manual start or restart option.

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

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

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

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

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

```
package IBMsdd is not installed
```

Configuring the Subsystem Device Driver

See *IBM TotalStorage Multipath Subsystem Device Driver User's Guide* for information about configuring and using the SDD.

Migrating with AIX 5L Version 5.2

The following sections describe the issues involved when you upgrade to, or remove, AIX 5L™ Version 5.2 with maintenance release 5200-01.

Storage unit migration issues when you upgrade to AIX 5L Version 5.2 maintenance release 5200-01

This topic discusses the migration issues when upgrading to AIX 5L version 5.2.

If a system has storage unit devices attached but the storage unit device specific object data manager (ODM) is not installed, the storage unit is configured as other FC disk. When you upgrade the system to AIX 5L Version 5.2 with the 5200-01 with the Recommended Maintenance package, the storage unit is no longer configured as other FC disk. These devices are configured as MPIIO other FC disk. A device is only configured as MPIIO other FC disk if the device is MPIIO-capable and no device-specific ODM is installed.

The following sections outline the differences between the other FC disk and MPIIO other FC disk configurations. A path is each physical connection between the host system and the device. A path control module is a device-specific module that manages device I/O across its paths.

A device configured as other FC disk has the following properties:

- It contains multiple device instances that are created for each path where the device was detected.
- It supports user-changeable device attributes.
- It can migrate to a vendor-specific device when vendor-supplied, device-specific ODM predefines are installed.
- It is meant to be in a transitory state during boot and install. Always install the vendor-specific device ODM predefines before using the device in a production environment.

A device configured as MPIIO other FC disk has the following properties:

- It contains only one device instance and multiple path instances. It also contains one path instance for each physical connection between the host system and the device.
- It supports user-changeable device attributes. There might be additional attributes that are specific to a path control module.
- It can migrate to a vendor-specific device when you install vendor-supplied device-specific ODM predefines.
- It presently is not supported by SDD path-management products. To support SDD path management, you must install the vendor-specific non-MPIO ODM predefines. Your attempt to control a device that is configured as an MPIIO device with SDD can produce undetermined results. Data integrity issues can exist if you operate the device in this configuration.
- It is supported in a production environment. Device-specific vendor ODM predefines are not required to be installed before you use the device in a production environment. A device-specific MPIIO-capable ODM is available for the storage unit.
- It allows you to install and boot to an MPIIO device.

If a system has storage unit devices attached and the storage unit device-specific ODM is installed, the storage unit is configured as storage unit-XXX disks. When the system is upgraded to AIX 5L version 5.2 with the 5200-01 Recommended Maintenance package, the storage unit does migrate the MPIIO other FC disk configuration. The MPIIO other FC disk configuration is only used for devices that support AIX MPIIO and do not have device-specific ODM installed.

MPIO-capable and non-MPIO versions of the storage unit ODM are available. The non-MPIO version is compatible with current releases of SDD. The MPIIO-capable version of the storage unit ODM is not compatible with current releases of SDD. For

more information on the Subsystem Device Driver (SDD), go to <http://www-1.ibm.com/servers/storage/support/software/sdd.html>.

For information about running MPIO on a pSeries host system, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Storage unit migration issues when you remove the AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package support

This section discusses the migration issues when removing AIX 5L version 5.2.

If you remove AIX 5L version 5.2 with the 5200-01 Recommended Maintenance package support after you configure devices as MPIO other FC disks, consider the following migration issues.

Migration back to the other FC disk configuration can occur if you remove the MPIO other FC disk support. In this case, where the update is uninstalled with the force option, the AIX version 5.2 release of MPIO manages the migration. If the system reboots, the device instance is in the defined state. During migration, the device instance is left in the defined state and a new other FC disk instance is created. If the system does not reboot and the device instance is in the defined state, the device instance remains in the defined state and a new other FC disk instance is created.

If the system does not reboot and the device instance is in the available state, the device instance remains unchanged. There might also be other FC device instances that are created. If the MPIO other FC device is not open, an other FC device instance is created for each path where the device is detected. If the MPIO other FC device instance is in the open state, no further other FC device instances are created. This occurs because the MPIO other FC device instance would have already issued a **SCIOSTART** command to the fibre-channel adapter for each of the paths. The fibre-channel adapter does not allow two devices with the same worldwide name and worldwide node name to exist in its internal data structures.

If other FC device instances are created, sending I/O to the device while it is configured as both MPIO other FC and other FC device might cause indeterminate device behavior or data corruption.

Recommendation: Reboot the system to clear this condition. Once you reboot the system, the MPIO other FC device instance is in the defined state and you can remove it with the **odmdelete** command. You cannot remove the device if you run the **rmdev** command due to missing predefines.

For information about running MPIO on a pSeries host system, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

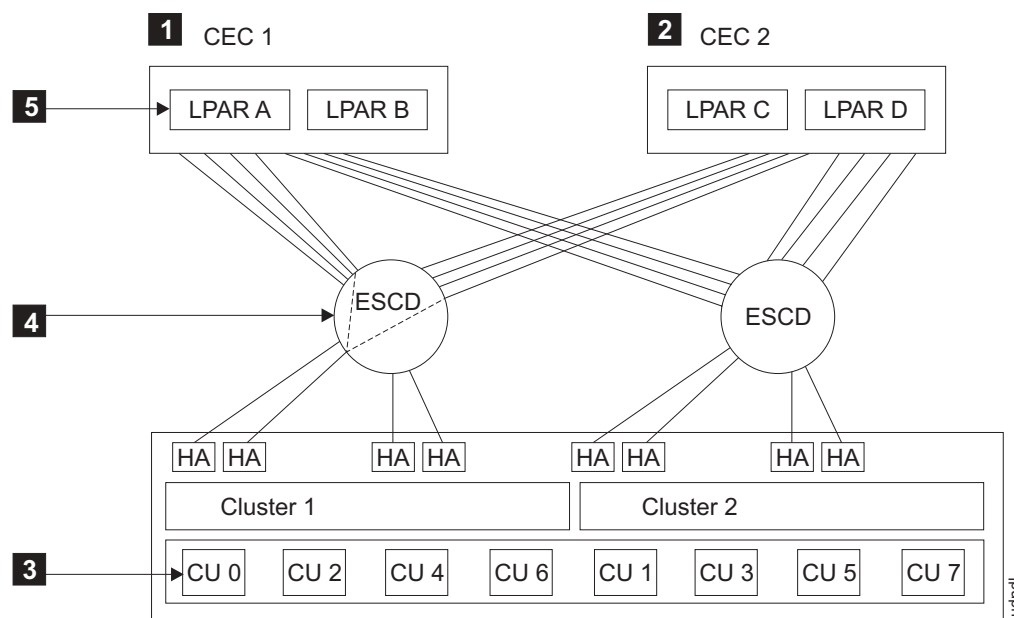
Chapter 10. Attaching to an IBM S/390 or IBM eServer zSeries host

This chapter describes the host system requirements to attach the IBM S/390 or zSeries host system to the storage unit with either an ESCON adapter or a FICON adapter.

ESCON connectivity on zSeries hosts

ESCON connectivity on zSeries hosts

Figure 48 shows how a storage unit that is attached through ESCON links to different computer-electronic complexes and logical partitions (LPARs). Figure 48 also shows a configuration that is designed for optimum availability. For optimum availability, make the ESCON host adapters available through all bays and spread the connections to any single host across both clusters. For optimum performance, have at least eight host adapter ports installed on four ESCON host adapter cards in the storage unit. This setup ensures the best performance from the attached systems.



Legend

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

Figure 48. ESCON connectivity

For information about how to use parallel access volumes (PAVs) for S/390 hosts and zSeries hosts, see the *IBM TotalStorage DS8000 User's Guide*. The guide

includes an overview of PAV and PAV requirements, and an example of the input/output configuration program (IOCP).

Attaching with ESCON adapters

The following section describes how to attach a storage unit with ESCON adapters.

Controller images and interconnections for ESCON adapters on zSeries hosts

This topic provides information about ESCON adapters on zSeries hosts.

For zSeries host attachment using ESCON, only the first 16 addresses of the LSSs can be used.

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

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

The controller images can handle the following ESCON interconnections:

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

Host adapters, cables, distances, and specifications for ESCON adapters on zSeries hosts

This section provides information for host adapters, cables, distances, and specifications for ESCON adapters on zSeries hosts.

Each ESCON host adapter connects to both clusters. A storage unit emulates 0 through 16 of the 3990 logical controllers. Half the logical controllers are in Cluster 1 and half in Cluster 2. Because the ESCON adapter connects to both clusters, each adapter can address 16 logical controllers.

Each ESCON host adapter provides two host connections. You must order two ESCON cables for each adapter for S/390 hosts and zSeries hosts.

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

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

An ESCON host channel can connect to more than one storage unit port through an ESCON director. The S/390 or zSeries host system attaches to one port of an ESCON host adapter in the storage unit. Each storage unit adapter card has two ports.

See the *IBM TotalStorage DS8000 Introduction and Planning Guide* for a list of the ESCON host adapter features codes. The publication also contains the number of ESCON host adapters, cable group number, number of cables, and connector IDs to order for the cables.

Cable lengths and types for ESCON adapters on zSeries hosts

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

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

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

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

Logical paths and path groups for ESCON adapters on zSeries hosts

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

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

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

Data transfer for ESCON adapters on zSeries hosts

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

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

Directors and channel extenders for ESCON adapters on zSeries hosts

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

Identifying the port for TSO commands on zSeries hosts

Figure 49 on page 102 helps you identify the port ID for the MVS TSO commands for S/390 and zSeries hosts. The numbers like 806, 804, 802, and 800 are internal microcode numbers. The numbers like 00, 04, 08, and 0C are tag numbers. You

need tag numbers for the path setup. To determine the tag number, use the **devserv** command with the **rcd** (read configuration data) parameter.

Slot#	Cluster 0				Cluster 1			
	Board 0 (CPI4)		Board 1 (CPI6)		Board 0 (CPI5)		Board 1 (CPI7)	
	1	2	3	4	1	2	3	4
AdapterID (8xx for NP and 9xx for PC)								
Port 0	806	804	802	800	816	814	812	810
Port 1	807	805	803	801	817	815	813	811
Tag Byte 1								
Port 0	00	04	08	0C	20	24	28	2C
Port 1	01	05	09	0D	21	25	29	2D
Logical Name (cpssnpXXX, cpsspcXX0)								
Port 0	400	410	420	430	600	610	620	630
Port 1	401	411	421	431	601	611	621	631

Figure 49. Port identification for S/390 and zSeries TSO commands

Migrating from ESCON to native FICON

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

Migrating from ESCON to native FICON on zSeries hosts: FICON support

FICON is supported on all 2107 models.

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

Migrating from ESCON to native FICON on zSeries hosts: Native ESCON configuration

This section provides a ESCON configuration example.

Figure 50 on page 103 shows an example of a native ESCON configuration. The configuration shows an S/390 or zSeries host that has four ESCON channels attached to the storage unit through two ESCON directors. The channels are grouped into a channel-path group for multipathing capability to the storage unit ESCON adapters.




Figure 50. Example of an ESCON configuration

Migrating from ESCON to native FICON on zSeries hosts: Mixed configuration

This section provides an ESCON migration example with S/390 or zSeries hosts.

Figure 51 on page 104 shows another example of a S/390 or zSeries host system with four ESCON channels. In this example, two FICON channels are added to an S/390 or zSeries host. The illustration also shows the channel path group and FICON directors through which the two FICON adapters are installed in the storage unit.

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

Figure 51 on page 104 also shows four ESCON adapters and two ESCON directors. This configuration gives you the most flexibility for future I/O changes.

Figure 51 on page 104 illustrates the FICON channels that have been added to the existing ESCON channel path group. Because the channel path group has ESCON and FICON channel paths, it makes migrating easier. This intermixing of types of channel-paths enables you to nondisruptively add FICON channels to the host and to add FICON adapters to the storage unit.

Note:

1. The configuration in Figure 51 on page 104 is supported for migration only.
2. Do not use this configuration for an extended period.

3. Migrate from a mixed channel-path group configuration to an all FICON channel-path group configuration.




Figure 51. Example of an ESCON configuration with added FICON channels

Migrating from ESCON to native FICON on zSeries hosts: FICON configuration

This section provides a FICON configuration example.

Figure 52 on page 105 illustrates how to remove the ESCON paths. The S/390 or zSeries host has four ESCON channels connected to two ESCON directors. The S/390 or zSeries host system also has two FICON channels.

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

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

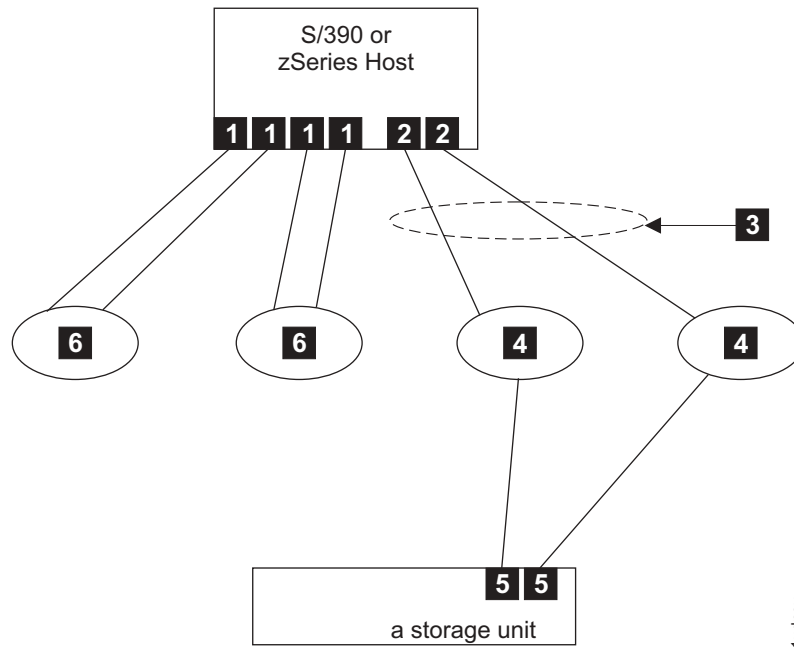


Figure 52. Example of a native FICON configuration with FICON channels that have been moved nondisruptively

Migrating from a FICON bridge to a native FICON attachment

This section shows how to migrate from a FICON bridge to a native FICON attachment.

Migrating from a FICON bridge to a native FICON attachment on zSeries hosts: FICON bridge overview

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

Migrating from a FICON bridge to a native FICON attachment on zSeries hosts: FICON bridge configuration

This section provides a FICON bridge example.

Figure 53 on page 106 shows an example of how to configure a FICON bridge. It also shows an S/390 or zSeries host with two FICON channels attached to two FICON bridges. You can attach the storage unit through the channel-path group to four ESCON links.

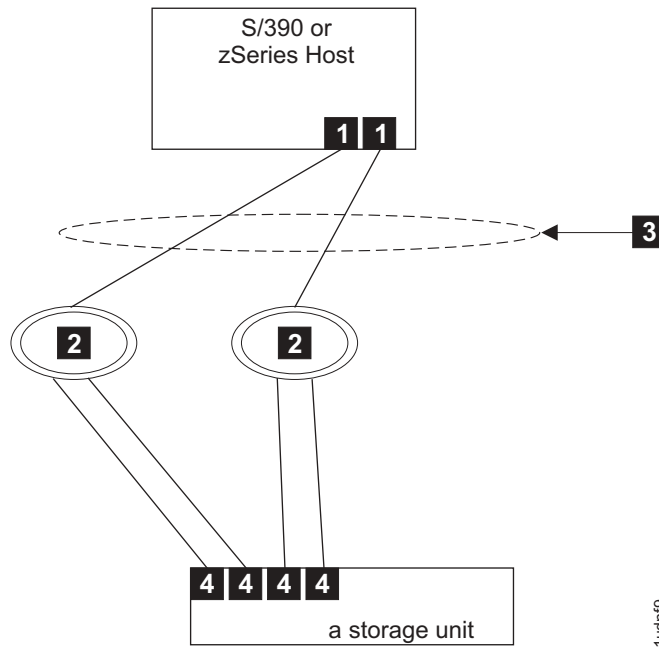


Figure 53. Example of how to configure a FICON bridge from an S/390 or zSeries host system to a storage unit

Migrating from a FICON bridge to a native FICON attachment on zSeries hosts: Mixed configuration

This section provides a FICON mixed configuration example.

Figure 54 on page 107 shows an example of an S/390 or zSeries host system with one FICON channel and one FICON director through a channel-path group and FICON host adapter to the storage unit. Figure 54 on page 107 also shows an S/390 or zSeries host system with one FICON channel and one ESCON director with a FICON bridge through a channel-path group and two ESCON adapters.

Figure 54 on page 107 shows that one FICON bridge was removed from the FICON configuration. The FICON channel that was connected to that bridge is reconnected to the new FICON director. The storage unit FICON adapter is connected to this new director. The channel-path group was changed to include the new FICON path. The channel-path group is a mixed ESCON and FICON path group. I/O operations continue to the storage unit devices across this mixed path group. Access to the storage unit devices is never interrupted because all the actions are nondisruptive.

Note:

1. The configuration in Figure 54 on page 107 is supported for migration only.
2. Do not use this configuration for an extended period.
3. Migrate from a mixed channel-path group configuration to an all FICON channel-path group configuration.




Figure 54. Example of how to add a FICON director and a FICON host adapter

Migrating from a FICON bridge to a native FICON attachment on zSeries hosts: Native FICON configuration

This section provides a native FICON configuration example.

Figure 55 on page 108 shows an S/390 or zSeries host system with two FICON channels connected to two FICON directors through a channel-path group to two FICON adapters. Note that the second bridge has been removed and a FICON director has been added. The channel-path group has only the FICON paths.




Figure 55. Example of the configuration after the FICON bridge is removed

Attaching with FICON adapters

This section tells you how to configure the storage unit for a FICON attachment.

Configuring the storage unit for FICON attachment on zSeries hosts

This topic provides configuration information for FICON attachment on zSeries hosts.

You can perform a FICON channel attachment on the IBM TotalStorage DS8000. See the *IBM TotalStorage DS8000 Introduction and Planning Guide* for a list of the FICON host adapter features codes. This publication also contains the number of FICON host adapters, cable group numbers, number of cables, and connector IDs to order for the cables.

Attachment considerations for attaching with FICON adapters on zSeries hosts

This section describes some things you should consider before you configure your system with a FICON interface.

Note: For a list of open-systems hosts, operating systems, adapters and switches that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Setting up ESCON and FICON links

If the system requires x ESCON links, where x is the number of links to get the performance and availability attributes you want, you must consider the number of FICON links you need. For example, you can map 4 ESCON links to a single FICON link and maintain approximately equivalent performance.

If the ESCON channel use is low, you can map 6 or 8 ESCON links to a single FICON link. If the ESCON channel use is high, you can map 2 ESCON links to a single FICON link.

Multipathing for ESCON and FICON

Consider the difference between the path groups when you compare FICON to ESCON. For example, for ESCON, you can configure 4 or 8 paths per path group from a host to a storage unit. For FICON, you can configure at least four paths in the path group to maximize performance.

Most ESCON controllers initiate channel command execution that partially synchronizes the lower disk drive interface with the upper channel interface. This channel command allows you only a very short time to reconnect. The consequence of limited time is a reconnection that can fail. Having eight paths in the path group, minimizes the number of missed reconnections. Increasing the number of path groups does not minimize the number of missed reconnections substantially. If you use eight paths in path groups, you can increase the overall throughput.

For FICON controllers, there is no synchronization between the lower disk drive interface and the upper channel interface. The number of paths in the path group depends on the throughput requirement. If it takes x paths to satisfy the throughput requirement, where x is the number of paths, set the path group to x .

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

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

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

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

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

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

Attaching to a FICON channel on a S/390 or zSeries host

This topic describes attaching to a FICON channel on a S/390 or zSeries host.

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

You can attach the FICON channels directly to a storage unit or you can attach the FICON channels to a fibre-channel switch. When you attach the FICON channels directly to a storage unit, the maximum number of FICON attachments is 16. Sixteen is the maximum number of host adapters you can configure in a storage unit. When you use a storage unit host adapter to attach to FICON channels either directly or through a switch, the adapter is dedicated to FICON attachment. It cannot be simultaneously attached to fibre-channel protocol hosts.

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

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

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

However, you can use cascaded switches on zSeries 800 and 900 model host systems with LIC patches installed and on all other zSeries models. To support cascaded switches, the fabric must be what is generally called a high integrity fabric. A high integrity fabric ensures data integrity by detecting any miscabling within a FICON cascaded switch fabric that could lead to I/O data being delivered to the wrong destination. A high integrity fabric is one that both supports and is configured to use fabric-binding and insistent domain IDs. Fabric-binding is the ability of the fabric to prevent another switch from being added operationally to the fabric without the new switch being properly planned for in the fabric and for the switch to be configured with the high integrity attributes. An insistent domain ID is a function of the fabric that does not allow the switch address to be automatically changed when a duplicate switch address is added to the fabric. Instead, fabrics that use insistent domain IDs require an operator's overt action to change a switch address. The customization of fabric-binding and the setting of an insistent domain ID are normally done only at switch installation time or reinstallation time.

Note: Not all FICON switches support high integrity fabrics. Such switches can only be used in a single-switch FICON fabric.

With fibre-channel and FICON-intermix mode, both fibre-channel protocol (FCP) and FICON upper-level protocols can be supported within the same director when deployed independently by port. (Director ports operate in either fibre-channel or FICON mode). The storage unit supports the fibre-channel and FICON-intermix mode on the following directors:

- CNT FC/9000 Fibre Channel Director (IBM 2042 Models 001, 128, and N16).
 - Implementation details and operational information for using intermix mode with the CNT FC/9000 is available at <http://www.inrange.com/ibm>.
 - A list of supported environments (servers, operating systems, adapters) and hardware and software prerequisites for the CNT FC/9000 is available at <http://www.storage.ibm.com/ibmsan/products/directors/index.html>.
- McDATA ED-6064 Enterprise Fibre Channel Director (IBM 2032 Model 001)
 - Implementation details and operational information for using intermix mode with the McDATA ED-6064 is available at <http://www.mcdata.com/ibm>.
 - A list of supported environments (servers, operating systems and adapters) and hardware and software prerequisites for the McDATA ED-6064 is available at <http://www.storage.ibm.com/ibmsan/products/2032/6064/>.
- Brocade 2109-M12 and 2109-M14 Fibre Channel Director
 - Implementation details and operational information for using intermix mode with the Brocade 2109-M12 and 2109-M14 is available at <http://www.brocade.com>.
 - A list of supported environments (servers, operating systems, adapters) and hardware and software prerequisites for the Brocade 2109-M12 and 2109-M14 is available at <http://www.storage.ibm.com/snetwork/index.html>.
- Cisco MDS 9509 and MDS 9506 Fibre Channel Directors, and MDS 9216 Fabric Switch
 - Implementation details and operational information for using intermix mode with the Cisco is available at <http://www.cisco.com/go/ibm/storage>.
 - A list of supported environments (servers, operating systems, adapters) and hardware and software prerequisites for the Cisco is available at <http://www.cisco.com/ibmsan/cisco/index.html>.

Registered state-change notifications (RSCNs) on zSeries hosts

This topic provides information about registered state-change notifications (RSCNs) on zSeries hosts.

McDATA and CNT switches ship without any configured zoning. This unzoned configuration is also referred to as enabling the default zone on some McDATA switches. This configuration enables all ports in the switch with a fibre-channel connection to communicate with each other and receive registered state change notifications (RSCNs) about each other. You can set the zones. Here are some recommendations:

- If you have FICON-only environments, do not disable the default zone.
- If you have FICON environments and plan to add (or already have) fibre-channel intermix on the switch, ensure that all FICON host adapters and all FICON-capable device ports are in an explicitly defined common zone.

- If you have open-systems traffic with fibre-channel protocol, continue to create zones that include the host ports and all desired storage ports. Multiple overlapping zones should be used in this case to support communication from multiple host ports to share device ports.

Linux for S/390 and zSeries

This section tells you how to attach a storage unit to an S/390 or zSeries host system that runs Linux.

Attaching a storage unit to an S/390 or zSeries host running Linux

This topic provides the references you need to attach a storage unit to an S/390 or zSeries host running Linux.

You can use either direct or switched attachment to attach a storage unit to an S/390 or zSeries host system that runs SUSE SLES 8 or 9 or Red Hat Enterprise Linux 3.0 with current maintenance updates for ESCON and FICON. For the most current list of supported hosts and release levels, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

For up-to-date information about the versions of Linux you can run on an S/390 or zSeries host, see <http://www.suse.de/us/business/products/server/>.

Disabling automatic system updates

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

If your host is attached to the DS8000 series and is running the IBM Multipath Subsystem Device Driver (SDD), you should consider turning off this automatic update feature because some drivers supplied by IBM, like SDD, are dependent on a specific kernel and cannot function in the presence of a new kernel. Similarly, host bus adapters drivers must be compiled against specific kernels to function optimally. By allowing automatic update of the kernel, you risk an unexpected impact to your host system.

Running Linux on an S/390 or zSeries host

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

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

- Native
Linux can run on the entire machine with no other operating system.
- Logical Partition (LPAR)
You can logically partition the hardware into a maximum of 15 separate LPARs per logical channel subsystem. For example, you can have S/390 applications in one partition, VM and VSE applications in another partition, and Linux applications in another partition.
- z/VM guest support

You can run Linux as a virtual machine using z/VM. VM provides virtualization of central processing units, I/O subsystems, and memory. You can have hundreds of Linux systems running on a single S/390 or zSeries host. For example, you can offer a complete Linux server environment to each of the application developers and host production systems all on the same S/390 or zSeries host.

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

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

Attaching FCP adapters on zSeries hosts

This section tells you how to attach fibre-channel protocol (FCP) adapters to a zSeries host running Linux.

Attaching FCP adapters on zSeries hosts running Linux

This section provides the information you need to attach FCP adapters on zSeries hosts running Linux.

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

Linux for S/390 (31-bit)

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

Linux for zSeries (64-bit)

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

FCP support

Linux for zSeries can access FCP and FICON controllers and devices with the appropriate I/O driver support. Linux can run either natively in a logical partition or as a guest operating system under z/VM Version 4 Release 4 or later. For the most current list of supported operating system releases see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

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

Accessing the storage unit using FCP for a zSeries host

This section lists the high level steps you must take to access the storage unit using FCP.

1. Gather the prerequisite information about the SAN setup. See “Setting up the storage unit for zSeries hosts running Linux” for a list of the information that you must collect.
2. Set up the storage unit and configure the LUNs for use with Linux for zSeries. See “Setting up the storage unit for zSeries hosts running Linux” for the steps you must perform to set up the storage unit.
3. Manually load the required modules and parameters to access the disks. See “Viewing the LUNs for zSeries hosts running Linux” on page 115 and “Setting up the Linux system to use the FCP device on zSeries hosts” on page 116 for the steps you must perform to view LUNs and set the Linux system to use FCP devices.
4. Make the disks permanently accessible. See “Adding additional permanent devices for zSeries hosts running Linux” on page 117 for the steps you must perform to add permanent devices.

Setting up the storage unit for zSeries hosts running Linux

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

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

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

You can obtain the data from the hardware management console (HMC), the IBM TotalStorage DS Storage Manager, and the SAN switch. Figure 56 shows examples of the prerequisite information you must have to run FCP Linux on the zSeries.

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

Figure 56. Example of prerequisite information for FCP Linux on zSeries

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

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

1. Logon to the IBM TotalStorage DS Storage Manager and click **Real-time manager**, click **Manage hardware**, click **Host systems**.
2. From the **Select action box**, click **Create**, then click **Go**.

3. From the **General host information** screen, complete the following fields for each fibre-channel host adapter. When finished, click **OK**.
 - Host type
 - Nickname
 - Description
4. From the **Define host ports** screen, specify the host ports for this host. Click **Add** to add each host port to the defined host ports table.
5. From the **Define host WWPN** screen, specify the world-wide port names for the selected hosts. When finished, click **Next**.
6. From the **Select storage images** screen, specify the storage image for attachment. Click **Add** to add the storage image to the selected storage images table. When finished, click **Next**.
7. From the **Specify storage image parameters** screen, specify the following parameters for each storage image.
 - Host attachment identifier
 - Volume Group for host attachment
 - Type of I/O port (any valid storage image I/O port or a specific I/O port)
8. Click **Apply assignment button** to apply the current attachment assignment. Use this button to go through this page repeatedly for each host attachment identifier that you want to assign to the storage image. When finished, click **OK**.
9. From the **Verification** screen, verify the data that represents the new host attachments. When finished, click **Finish**.

Viewing the LUNs for zSeries hosts running Linux

This topic provides instructions for viewing the LUNs for zSeries hosts running Linux.

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

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

You can obtain the data from the hardware management console (HMC), the IBM TotalStorage DS Storage Manager, and the SAN switch. Figure 57 shows examples of the prerequisite information you must have to run FCP Linux on the zSeries.

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

Figure 57. Example of prerequisite information for FCP Linux on zSeries

To view the LUNS of the volumes assigned to your Linux system, use the IBM TotalStorage DS Storage Manager.

Setting up the Linux system to use the FCP device on zSeries hosts

This topic provides instructions for setting up the Linux system to use the FCP device on zSeries hosts.

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

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

You can obtain the data from the hardware management console (HMC), the IBM TotalStorage DS Storage Manager, and the SAN switch. Figure 58 shows examples of the prerequisite information you must have to run FCP Linux on the zSeries.

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

Figure 58. Example of prerequisite information for FCP Linux on zSeries

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

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

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

Adding additional permanent devices for zSeries hosts running Linux

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

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

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

You can obtain the data from the hardware management console (HMC), the IBM TotalStorage DS Storage Manager, and the SAN switch. Figure 59 shows examples of the prerequisite information that you must have to run FCP Linux on the zSeries.

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

Figure 59. Example of prerequisite information for FCP Linux on zSeries

1. Choose one of the following methods to add devices: write a script or add the device manually.
 - a. To add more than one device to your SCSI configuration, write a small script with all the parameters included. Figure 60 shows an example of such a script. This is an optional step.

```
rmmod zfc
modprobe qdio
modprobe scsi_mod
insmod zfc map="\
0x0600 1:0x5005076300c89589 0:0x5301000000000000;\
0x0600 1:0x5005076300c89589 1:0x5302000000000000;\
0x0600 1:0x5005076300c89589 2:0x5300000000000000"
modprobe sd_mod
```

Figure 60. Example of a script to add more than one device

- b. Alternatively, you can also add SCSI devices to an existing configuration with the **add_map** command. After using this command, you must manually make the devices known to the SCSI stack. Figure 61 on page 118 shows an example of how to add SCSI devices through the **add_map** command.

```
linuxa:~#echo "0x0600 0x00000001:0x5005076300c89589 \
0x00000001:0x5302000000000000">/proc/scsi/zfcp/add_map
linuxa:~#echo "scsi add-single-device 0 0 1 1">/proc/scsi/scsi
```

Figure 61. Example of how to add SCSI devices through the **add_map** command

2. To make the devices available permanently, after a reboot, you must create a new initial ramdisk (RAM that is accessed like a disk) that contains the necessary modules and parameter information. First, save the module parameters in the configuration file `/etc/zfcp.conf`, as shown in Figure 62.

```
linuxa:~#cat /proc/scsi/zfcp/map >/etc/zfcp.conf
linuxa:~#cat /etc/zfcp.conf
0x0600 0x00000001:0x5005076300c89589 0x00000000:0x5301000000000000
0x0600 0x00000001:0x5005076300c89589 0x00000001:0x5302000000000000
0x0600 0x00000001:0x5005076300c89589 0x00000002:0x5300000000000000
```

Figure 62. Saving the module parameters in the `/etc/zfcp.conf` file

3. Next, create a new ramdisk with the **mk_initrd** utility and then run the **zipl** utility to update the initial program load record to point to the new ramdisk.
4. Enable multipath support, which is provided by the logical volume manager kernel for Linux z/FCP.

- a. Create physical volumes with the following commands:

```
pvccreate /dev/sda1
pvccreate /dev/sdb1...
```

- b. Put the physical volumes into a volume group with the following command:

```
vgcreate fcpvg /dev/sda1 /dev/sdb1...
```

- c. Create as many logical volumes as you want with the following command:

```
lvcreate --size 16G fcpvg
```

- d. Enable the alternate paths to the physical volumes with the **pvpath** command.

```
pvpath --path0 --enable y /dev/sda1
pvpath --path1 --enable y /dev/sda1
```

If both paths have a weight of 0, they will loadbalance.

5. These configurations will yield the SCSI device `/dev/sda` - `/dev/sdc` accessed on the first path and the SCSI device `/dev/sdd` - `/dev/sdf` accessed on the second path, as shown in Figure 63.

```
rmmod zfcp modprobe qdio modprobe scsi_mod insmod zfcp map="\ 0x0600 1:0x5005076300c89589
0:0x5301000000000000;\ 0x0600 1:0x5005076300c89589 1:0x5302000000000000;\
0x0600 1:0x5005076300c89589 2:0x5300000000000000;\ 0x0700 1:0x5005076300ca9589
0:0x5301000000000000;\ 0x0700 1:0x5005076300ca9589 1:0x5302000000000000;\
0x0700 1:0x5005076300ca9589 2:0x5300000000000000" modprobe sd_mod
```

Figure 63. Example of Logical Volume Manager Multipathing

Chapter 11. Attaching to an Intel host running Linux

This chapter provides instructions, requirements, and considerations for attaching an Intel[™] host system running Linux to a storage unit.

Supported adapter cards for an Intel host running Linux

This topic lists the supported adapter cards for attaching to an Intel host running Linux.

This section describes how to attach an Intel server running Linux to a storage unit with the following fibre-channel adapters:

- QLogic QLA2310F adapter card
- QLogic QLA2310FL adapter card
- QLogic QLA2340FL adapter card
- QLogic QLA2342FL adapter card
- Emulex LP9002DC adapter card
- Emulex LP9002L adapter card
- Emulex LP9402DC adapter card
- Emulex LP9802 adapter card
- Emulex LP10000 adapter card
- Emulex LP10000DC adapter card
- Netfinity adapter card

For up-to-date information about the servers, operating systems, distribution levels, host adapters, and fabric connections that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Note: The steps to install and configure adapter cards are examples. Your configuration might be different.

Attaching with fibre-channel adapters

You can attach an Intel server that is running either the Red Hat Enterprise Linux, SUSE SLES, Red Flag Linux Advanced Server, and Asianux versions of the Linux operating system to an IBM storage unit.

For information about the most current version of the kernel and the switches that are supported, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Attachment requirements for an Intel host running Linux

This section provides attachment requirements for an Intel host running Linux

This section lists the requirements for attaching the storage unit to your host system:

- Check the LUN limitations for your host system.
- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *User's Guide* is on the TotalStorage CD that you receive with the storage unit.

- For details about the release level for your operating system, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Either you or an IBM service support representative must perform the following tasks to install and configure a storage unit:

1. Install the storage unit.
2. Define the fibre-channel host system with the worldwide port name identifiers. For the list of worldwide port names, see “Locating the worldwide port name (WWPN),” on page 225.
3. Define the fibre-port configuration if you did not do it when you installed the storage unit or fibre-channel adapters.
4. Configure the host system for the storage unit by using the instructions in your host system publications.

Attachment considerations for an Intel host running Linux

This section provides attachment considerations for an Intel host running Linux.

For the most current Linux versions, servers, and switches, see <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

See <http://www.redhat.com> or <http://www.suse.com> to download the vendor kernel.

The DS8000 supports eServer BladeCenter HS20 with blades 8678, 8832, and 8839 when operating in a Red Hat Enterprise Linux 3.0 or 4.0 or SUSE SLES 9, 32-bit environment. See <http://www-1.ibm.com/servers/storage/disk/ds8000> for more information about supporting BladeCenter HS20.

The storage unit supports both direct and fabric attachment to the eServer BladeCenter™ JS20 when operating with AIX 5.2H, or Red Hat Enterprise Linux 3.0 or 4.0, or Linux SLES 9. See the <http://www.ibm.com/servers/storage/disk/ds8000/interop.html> for more information about supporting BladeCenter JS20.

Disabling automatic system updates

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

If your host is attached to the DS8000 series and is running the IBM Multipath Subsystem Device Driver (SDD), you should consider turning off this automatic update feature because some drivers supplied by IBM, like SDD, are dependent on a specific kernel and cannot function in the presence of a new kernel. Similarly, host bus adapters drivers must be compiled against specific kernels to function optimally. By allowing automatic update of the kernel, you risk an unexpected impact to your host system.

Installing the Emulex adapter card for an Intel host running Linux

This section provides instructions for installing the Emulex adapter card for an Intel host running Linux.

This section tells you how to attach a storage unit to a Linux host system with an Emulex adapter card.

Note: The LP9402DC adapter card is a 3.3 volt card only and requires a 3.3 volt PCI slot.

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

- Target
 - Public initiator
 - Private initiator
 - Target and public initiator
 - Target and private initiator
1. Record the Institute of Electrical and Electronics Engineers (IEEE) number that is printed on the card. You can use the IEEE number to determine the WWPN.

Note: You can find the WWPN of your adapter cards by issuing the `cat/proc/scsi/adapter_type/n` command, where *adapter_type* is the host adapter type and *n* is the host adapter number for your card. Your host adapter type can be either `qlxxxx` for QLogic adapters or `lpfc` for Emulex adapters.

2. Refer to the installation instructions provided by your host adapter vendor for their specific installation instructions. Go to <http://www.emulex.com>

Downloading the Emulex adapter driver for an Intel host running Linux

This section provides instruction for downloading the Emulex adapter driver for an Intel host running Linux.

1. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
2. Click **Interoperability matrix**.
3. Click **DS8000 interoperability matrix**.
4. Find the section for the current version of the driver and firmware that you want.
5. Go to <http://www.emulex.com>
6. Click **Driver, downloads, and documentation** from the left navigation pane. Click **OEM software and documentation**.
7. Select **IBM** from the drop down vendor list.
8. Click the link for the adapter to get the firmware, driver, and documentation that you need to install and download the adapter driver.

Installing the Emulex adapter driver for an Intel host running Linux

This section provides instructions for installing the Emulex adapter drivers for an Intel host running Linux.

You must be able to reference the readme documentation that comes with the host adapter.

The Emulex Web site contains detailed instructions to install, build, and load their adapter drivers. Refer to the readme documentation that comes with your host adapter for these detailed instructions that are unique to your host adapter. You must follow these instructions for a successful installation of the adapter driver. See <http://www.emulex.com> for the latest documentation.

Installing the QLogic adapter card for an Intel host running Linux

This section tells you how to attach an Intel host system running Linux to a storage unit with the QLogic QLA23xx adapter card.

Single- and dual-port fibre-channel interfaces with the QLogic adapter cards support the following public and private loop modes:

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

Note: The following steps are an example configuration. The configuration for your adapter might differ.

1. Install the QLogic QLA23xx adapter card in the host system.
2. Connect the cable to the storage unit port.
3. Restart the server.
4. Press Ctrl+Q to get to the FAST!Util menu panel.
5. Use FAST!Util to set the parameters to the values listed in the following table:

Table 8. Recommended settings for the QLogic adapter card for an Intel host running Linux

Parameters	Recommended settings
Basic Input/Output System (BIOS)	Disabled
Frame size	2048
Loop reset delay	5 (minimum)
Adapter hard loop ID	Disabled
Hard loop ID	0
Execution throttle	100
Fast command posting	Enabled
>4 GB addressing	Disabled (for 32-bit systems)
LUNs per target	0 or 128
Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes
Login retry count	<ul style="list-style-type: none">• 60 if you are not running the SDD• 8 if you are running the SDD
Port down retry count	<ul style="list-style-type: none">• 60 if you are not running the SDD• 8 if you are running the SDD
Driver load RISC code	Enabled
Enable database updates	No
Disable database load	No
IOCB allocation	256
Extended error logging	Disabled (You might want to enable this parameter if you use debugging.)

Table 8. Recommended settings for the QLogic adapter card for an Intel host running Linux (continued)

Parameters	Recommended settings
Type of connection	<ul style="list-style-type: none"> • Loop only • Point-to-point • Loop preferred (If you cannot use arbitrated loop, then default to point-to-point.) • Point-to point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop.) <p>Note: If you connect the storage unit directly to the host system, the option you select must match the port connections on the storage unit.</p>

6. Restart the server.

Downloading the current QLogic adapter driver for an Intel host running Linux

This section provides instructions for downloading the current QLogic adapter driver for an Intel host running Linux.

1. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
2. Click **Interoperability matrix**.
3. Click **DS8000 interoperability matrix**.
4. Find the section for the current version of the driver and firmware and driver you want.
5. Go to <http://www.qlogic.com>.
6. Click **Drivers / Software** from the **Quick Links** list.
7. Choose **IBM** from the **OEM and ISV approved Drivers and Firmware** section.
8. Navigate to the part of the page for the adapter you want to download, then click **Download**.

Installing the QLogic adapter driver for an Intel host running Linux

This section provides instructions for installing the QLogic adapter drivers for an Intel host running Linux.

You must be able to reference the readme documentation that comes with the host adapter.

The QLogic Web site contains detailed instructions to install, build, and load their adapter drivers. Refer to the readme documentation that comes with your host adapter for these detailed instructions that are unique to your host adapter. You must follow these instructions for a successful installation of the adapter driver. See <http://www.qlogic.com> for the latest documentation.

Defining the number of disk devices on Linux

This section provides information about defining the number of disk devices on Linux.

The maximum number of configurable devices that are available on supported Linux host systems is either 128, 256, or 2304. The standard Linux kernel uses a static major and minor number address mechanism. A special device file represents each disk device. For each default, there is a maximum of 16 partitions per disk. The major and minor numbers are each 8-bits.

Note: To determine what level of Linux is installed, use the following commands:

```
uname -a
cat /etc/*release*
```

Fibre-channel attached devices are managed as SCSI devices. Traditionally, there have been eight major numbers reserved for SCSI devices. Those major numbers are 8, 65, 66, 67, 68, 79, 70, and 71. For kernel levels that support 256 devices, the additional major numbers are 128, 129, 130, 131, 132, 133, 134, and 135. Newer kernels are moving to a dynamic configuration method that will use any available major numbers. The latest SUSE SLES 9 kernel errata include this change and impose a limit of no more than 144 major numbers allocated to SCSI disk devices.

There are 256 minor numbers available for each of the eight major numbers. The following formula provides the maximum number of devices for Linux host systems: Number of devices = (number of major numbers) x (number of minor numbers) ÷ (number of partitions). For example, Number of devices = 8 x 256 ÷ 16 = 128 or Number of devices = 144 x 256 ÷ 16 = 2304.

To set the maximum number of disk devices under Linux to properly detect all volumes, you need to set the `max_scsi_luns` option.

1. Edit the `/etc/modules.conf` file.
2. Add the following line: `options scsi_mod max_scsi_luns=128`.

Note: If you already have a `max_scsi_lun` option in the `modules.conf` file, ensure that the values is set to 128.

3. Save the file.
4. Run `mkinitrd` in order to rebuild the ram-disk associated with the current kernel. You can use the following command as an example of what `mkinitrd` command to run:

```
mkinitrd /boot/<image> <kernel>
```

5. Reboot the host.

If you load a host adapter device as a module in Linux, there is a kernel limit to the total number of disks that you can configure. This limit might be less than the total supported by the kernel (usually 128 or 256). The first module that is loaded on the system might configure discovered disks up to the maximum number of disks supported by the kernel. Subsequent drivers are limited to configuring a fixed, smaller number of drives. These drivers all share a pool of device structures that are statically allocated after the first host adapter driver is loaded. The number of allocated device structures (usually 40) is set with the `CONFIG_SD_EXTRA_DEVS` parameter when the kernel compiles. Traditionally, this value has been 40, but newer Red Hat kernels have it set to 128 or 256. On SUSE SLES 9 the dynamic configuration system eliminates this limit.

For example, consider a system with an internal SCSI Adaptec host adapter and one disk. The system contains fibre-connected disks and a `CONFIG_SD_EXTRA_DEVS` parameter set to 40. The Adaptec host adapter driver loads first and configures the single attached disk. At this point, the system is limited to a total of 41 SCSI disks instead of the full 128 or 256. If there are more than 40 disks visible in the fibre-channel fabric, the system configures and uses only the first 40. All of the fabric devices might be listed in the `/proc/scsi/scsi` directory, but only the first 40 are configured with the SCSI disk driver. You are more likely to encounter this constraint in multipath fabric situations.

Configuring the storage unit

This section provides instructions for configuring the storage unit.

Configuring the storage unit for an Intel host running Linux

Configuring the storage unit for an Intel host running Linux

Each of the attached storage unit LUNs has a special device file in the Linux directory `/dev`. There are maximums of 128 or 256 fibre-channel disks that are based on the major numbers that are available. For Red Hat, the operating system automatically adds the entries for all 128 devices. For SUSE, there are only special device files for the first 16 disks. You must create the device files for additional disks by using the **mknod** command.

The range of devices goes from `/dev/sda` (LUN 0) to `/dev/sddx` (LUN 127). Figure 64 shows an example of the range for the devices.

```
# ls -l /dev/sda
brw-rw---- 1 root disk 8, 0 Aug 24 2000 /dev/sda
```

Figure 64. Example of range of devices for a Linux host

Partitioning storage unit disks for an Intel host running Linux

Partitioning storage unit disks for an Intel host running Linux

Before you create a file system, partition the disk by using the `fdisk` utility. You have to specify the special device file of the disk you want to partition when executing `fdisk`. Figure 65 on page 126 shows an example of the different options for the `fdisk` utility.

```
# fdisk /dev/sdb

Command (m for help): m
Command action
a toggle a bootable flag
b edit bsd disklabel
c toggle the dos compatibility flag
d delete a partition
l list known partition types
m print this menu
n add a new partition
o create a new empty DOS partition table
p print the partition table
q quit without saving changes
s create a new empty Sun disklabel
t change a partitions system id
u change display/entry units
v verify the partition table
w write table to disk and exit
x extra functionality (experts only)
```

Figure 65. Example of different options for the fdisk utility

Figure 66 shows an example of a primary partition on the disk /dev/sdb.

```
Command (m for help): n

Command action
e extended
p primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-953, default 1): Enter
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-953, default 953): Enter
Using default value 953

Command (m for help): p

Disk /dev/sdb: 64 heads, 32 sectors, 953 cylinders
Units = cylinders of 2048 * 512 bytes

Device Boot Start End Blocks Id System
/dev/sdb1 1 953 975856 83 Linux
```

Figure 66. Example of a primary partition on the disk /dev/sdb

Assigning the system ID to the partition for an Intel host running Linux

Assigning the system ID to the partition for an Intel host running Linux

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

Figure 67 on page 127 shows the assignment of the Linux system ID to the partition (hex code 83).

```

Command (m for help): t
Partition number (1-4): 1

Hex code (type L to list codes): 83

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
SCSI device sdb: hdwr sector= 512 bytes. Sectors= 1953152 [953 MB] [1.0 GB]
sdb: sdb1
SCSI device sdb: hdwr sector= 512 bytes. Sectors= 1953152 [953 MB] [1.0 GB]
sdb: sdb1

WARNING: If you have created or modified any DOS 6.x partitions, please see the
fdisk manual page for additional information.
Syncing disks.
[root@yahoo /data]#

```

Figure 67. Example of assigning a Linux system ID to the partition

Creating and using file systems on the storage unit for an Intel host running Linux

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

After you partition the disk as described in “Partitioning storage unit disks for an Intel host running Linux” on page 125, the next step is to create a file system. Figure 68 shows an example of the EXT2 Linux file system (which is nonjournaled) by using the **mke2fs**.

```

[root@yahoo /data]# mke2fs /dev/sdb1
mke2fs 1.18, 11-Nov-1999 for EXT2 FS 0.5b, 95/08/09
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
122112 inodes, 243964 blocks
12198 blocks (5.00%) reserved for the super user
First data block=0
8 block groups
32768 blocks per group, 32768 fragments per group
15264 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376

Writing inode tables: done
Writing superblocks and filesystem accounting information: done
[root@yahoo /data]#

```

Figure 68. Example of creating a file with the mke2fs command

Figure 69 on page 128 shows an example of the EXT2 Linux file system (which is nonjournaled) by using the **mkfs** command.

```
[root@yahoo /data]# mkfs -t ext2 /dev/sdb1
mke2fs 1.18, 11-Nov-1999 for EXT2 FS 0.5b, 95/08/09
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
122112 inodes, 243964 blocks
12198 blocks (5.00%) reserved for the super user
First data block=0
8 block groups
32768 blocks per group, 32768 fragments per group
15264 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376
Writing inode tables: done
Writing superblocks and filesystem accounting information: done
[root@yahoo /data]#
```

Figure 69. Example of creating a file with the mkfs command

SCSI disk considerations for an Intel host running Linux

This section provides information for dynamically adding and removing SCSI disks.

Manually adding and removing SCSI disks

This section provides the steps you must follow to manually add and remove SCSI disks.

You can unconfigure and remove an unused SCSI disk with the following command:

```
echo "scsi remove-single-device H B T L" > /proc/scsi/scsi
```

where H B T L are the host, bus, target, and LUN IDs for the device, respectively.

When you add disks to the fabric, you can use either of the following two ways to force the Linux host to configure them:

- If the host adapter driver that controls the LUN is not in use, unload it and load the module again.
- If the driver cannot be unloaded and loaded again, and you know the host, bus, target and LUN IDs for the new devices, you can add them through the /proc/scsi/scsi file using the following command:

```
echo "scsi add-single-device H B T L" > /proc/scsi/scsi
```

where H B T L are the host, bus, target, and LUN IDs for the device, respectively.

Note: The Linux kernel does not assign permanent names for the fabric devices in the /dev directory. For example, a LUN might be /dev/sda. After a driver reload, the same LUN might become /dev/sdce. A fabric reconfiguration might also result in a shift in the host, bus, target and LUN IDs, which makes it unreliable to add specific devices through the /proc/scsi/scsi file.

LUN identification for the Linux host system

This section describes LUN identification for the Linux host system.

Some Linux drivers do not automatically configure any LUNs other than LUN 0. You might have to manually configure the other LUNs, depending on the Linux driver that you are using. You can find the list of assigned LUNs in the `/proc/scsi/scsi` file. Figure 70 shows an example of the `proc/scsi/scsi` file for a Linux host that only configures the first LUN, LUN 0, on each host adapter port.

```
# cat /proc/scsi/scsi
...
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM-PSG Model: DPSS-318350M F Rev: S9HA
  Type:   Direct-Access                      ANSI SCSI revision: 03
Host: scsi0 Channel: 00 Id: 15 Lun: 00
  Vendor: IBM Model: TP4.6 V41b3           Rev: 4.1b
  Type:   Processor                         ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM Model: 2105800               Rev: .294
  Type:   Direct-Access                     ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM Model: 2105800               Rev: .294
  Type:   Direct-Access                     ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM Model: 2105800               Rev: .294
  Type:   Direct-Access                     ANSI SCSI revision: 03
```

Figure 70. Example output for a Linux host that only configures LUN 0

You can create a script to configure every LUN that is assigned to the system. Use the **scsi add-single-device** command to consecutively configure all of the LUNs that are assigned to your host system. Write a script that repeats the **scsi add-single-device** command for each LUN on each ID for each host adapter. The script must scan all host adapter ports and identify all of the LUNs that are assigned to each port.

After you run the script, you can view all of the assigned LUNs in the `proc/scsi/scsi` file. Figure 71 shows an excerpt of an example `proc/scsi/scsi` file for a Linux host after a script has configured every LUN.

```
# cat /proc/scsi/scsi
...
Host: scsi3 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM Model: 2105800               Rev: .294
  Type:   Direct-Access                     ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM Model: 2105800               Rev: .294
  Type:   Direct-Access                     ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM Model: 2105800               Rev: .294
  Type:   Direct-Access                     ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 02
  Vendor: IBM Model: 2105800               Rev: .294
  Type:   Direct-Access                     ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 03
  Vendor: IBM Model: 2105800               Rev: .294
  Type:   Direct-Access                     ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 04
  Vendor: IBM Model: 2105800               Rev: .294
  Type:   Direct-Access                     ANSI SCSI revision: 03
...
```

Figure 71. Example output for a Linux host with configured LUNs

SCSI disk problem identification and resolution

This section describes a few problems that might occur on your storage unit when you are using SCSI disks.

- In response to errors, the kernel might permanently disable a LUN and log a message stating “device set offline” and the specific device. If this happens, there is no way to bring the LUN online. It might be possible to remove the LUN and add it back through the `/proc/scsi/scsi` directory or reload the host adapter driver. If those fail, it is necessary to reboot the host.
- The system might periodically list processes in the D-state (see the `ps` command help page), which corresponds to an uninterruptible sleep. In error situations, a process might become permanently stuck in this state, and require a system reboot to recover.
- The Linux kernel buffer cache is designed to discard dirty buffers after an input/output (I/O) error when the system memory resources are constrained. An application that is attempting to use the `fsync()` command to verify that its writes have completed successfully will receive an indication of success from the command once the writes successfully complete. Some kernels have a bug in the `kswapd` daemon, that makes it likely that the system will perceive itself to be in a state of constrained memory. Multipathing can reduce the risk of this silent data loss by providing a means to retry failed I/O operations and hide the failure from the buffer cache.

Support for fibre-channel boot

The following sections contain support information for fibre-channel boot.

Creating a modules disk for SUSE LINUX Enterprise Server 9.0

Use this process to create a modules disk for SUSE LINUX Enterprise Server 9.0.

1. Boot the server using the SLES 9 SP1 CD.
2. When you are prompted, select **Rescue System**.

Note: If you are using the x360/x440 servers, append “apic” to the kernel boot options string at the bottom of the menu screen.

3. Insert the SLES 9 CD 1 CD when you are prompted.
4. After the system boots, the Rescue login prompt appears. Log in as “root” and do not enter a password.
5. Remove the SLES 9 CD 1 CD, and insert the SLES 9 SP1 CD.
6. If you are using IDE CD-ROM drives, type `mount /dev/hda /cdrom` at the command line. If you are using SCSI, type `/dev/sda` instead of `hda`, or `/dev/sdb`, and so forth.
7. Change to the CD-ROM boot directory by typing `cd /cdrom/boot`.
8. Insert a diskette into the diskette drive. Create the modules disk by typing `dd if=modules2 of=/dev/fd0u1440`. When this is completed, a message is displayed that says “2880+0 records in 2880+0 records out”.
9. Reboot with the SLES 9 SP1 CD in the CD-ROM drive.

Installing Linux over the SAN without an IBM Subsystem Device Driver

This topic provides instructions for installing Linux over the SAN without an IBM Subsystem Device Driver.

Use this process to install Linux over the SAN without an IBM Subsystem Device Driver. These instructions only apply to Qlogic adapter cards or to the eServer BladeCenter HS20. These instructions do not work if you use them on a server running Emulex adapter cards.

1. Boot the server.
2. Press the <CTL-Q> message for the Fast!Util when it appears.
3. In the QLogic FAST!Util, select the controller that has your remote boot disks attached.
4. Click **Configuration Settings** → **Host Adapter Settings** and change the BIOS enabled setting to **Enabled**.
5. Click **Configuration Settings** → **Selectable Boot Settings**.
 - a. Change the Selectable boot setting to **Enabled**.
 - b. Set **(Primary) Boot Port Name,Lun:** to the first device in your list.

Note: If you have assigned multiple LUNs and are using any of them to boot from (that is, if you have file systems such as /boot, /, /usr, /var, or the swap partition), you must list the LUNs in the subsequent entries, as well.

6. Save your changes.
7. Back out of the FAST!Util and reboot.
8. Insert modules disk modules2 into the diskette drive (see the previous procedure for creating modules disks).
9. Put the SP1 CD in the CD-ROM drive and reboot.
10. Select **Manual Installation**.

Note: If you are using a Summit machine, you must add “apic” to the boot options at the bottom of the menu screen.

11. Select the appropriate keyboard and press Enter.
12. After manual installation, select the **Kernel modules (hardware drivers)** and **Load IDE/RAID/SCSI modules** options.
13. Select the QLA23xxxx module from the list, press Enter twice, leaving the parameters line blank.
14. Wait for the module to return from loading, and then press Enter.
15. Press the right-arrow key once to select the “Back” button, and press Enter again.
16. Select **Start Installation/Update** → **Install** → **CD-Rom**.
17. When you are prompted with “Please make sure that CD number 1 is in your drive!”, place the SLES 9 Install CD 1 CD into the CD-ROM drive and press Enter.
18. When you are prompted with “SUSE’s End User License for SLES”, select **Accept**, and then select the appropriate keyboard. Click **Accept**.
19. During detection, if you are asked whether you want to make a New Installation, select **New Installation** and continue.
20. Create the partitions. (Normally, there are three: one of approximately 200 MB for /boot, a second of approximately 2 GB for swap, and the rest for /.)
21. Make the package selections, making sure to include the kernel source package.
22. Check your boot device parameters. In most cases, add “apic=off” to the load string. If you have loaded the apic boot option at the outset, it is listed automatically. If you used an APIC-enabled installation and the apic boot

option is not listed automatically, add it. Failure to add the apic boot option results in a panic condition for the server at reboot.

23. Check your time zone parameter and change it to the appropriate time zone.
24. Check the clock setting. If it defaults to “UTC”, change it to “Local”.
25. Click **Next** and then click **Yes** when you are prompted. The system formats the remote disk, and then installs the packages.
26. Switch CDs when you are prompted.

Note: To boot Linux, follow the instructions above to allow the server to boot.

Updating to a more recent module without an IBM Subsystem Device Driver

Use this process to update to a more recent module without IBM Subsystem Device Driver.

1. You must boot into SUSE SLES 9 before you install the QLogic driver source and compile the necessary kernel and driver.
2. Download the current QLogic adapter driver. For additional information, see “Downloading the current QLogic adapter driver for an Intel host running Linux” on page 123.
3. Type `cat /etc/sysconfig/kernel` at the command prompt and make sure that the QLA23xx driver is listed.
4. Type `mk_initrd` to pick up the new driver in the ramdisk image for this kernel, and then reboot.

Note:

- a. If you are using lilo, you must rerun lilo to install the changed boot parameters into the MBR. To do this, type `mk_initrd && lilo` from the command line. GRUB users do not need to make any changes after running `mk_initrd`.
- b. If you ever need to upgrade the kernel, repeat this procedure for updating the module before you reboot. Make sure the `mk_initrd` command runs successfully before you reboot.

Installing Linux over the SAN with an IBM Subsystem Device Driver

Use this process to install Linux over the SAN with an IBM Subsystem Device Driver.

This procedure assumes that you have correctly configured the bootloader to boot from the single-pathed SAN device. It also assumes that the SDD rpm is installed on your system. This procedure describes how to copy SDD files into the initial ramdisk (`initrd`) and edit the `linuxrc` script, which is processed when the kernel mounts the initial ramdisk at boot time. For current information about SDD rpm, see the *IBM TotalStorage Multipath Subsystem Device Driver User's Guide*.

Perform the following steps to install Red Hat and SUSE with an IBM Subsystem Device Driver:

1. Make a backup of the existing `_initrd_file_` by typing the following commands:

```
cd /boot
cp _initrd_file_ _initrd_file.bak
```

2. Uncompress the image by typing the following command:

```
zcat _initrd_file_.bak > _initrd_file_
```

3. Set up a loopback device for the image by typing the following command:

```
losetup /dev/loop0 /path/to/your/_initrd_file_
```

4. Fix any errors that might exist on the file system by typing the following command:

```
e2fsck -f /dev/loop0
```

5. Determine the size of the initrd file system by typing the following command:

```
df /dev/loop0
```

Ensure that you have sufficient space in the /boot directory (or other home of your initrd files) to store considerably larger initrd files (for example, files with a size of 32 MB each). If there is not sufficient space in your /boot directory, you can perform the following steps in a temporary directory and then copy the compressed initrd file (a few megabytes instead of 32 MB) back into the /boot directory.

If the file system is not 32 MB or larger or does not have much free space, you must enlarge it to 32 MB by typing the following command:

```
losetup -d /dev/loop0  
dd if=/dev/zero of=_initrd_file_ seek=33554432 count=1bs=1
```

Note: If the file is already 32 MB or larger, do not perform this step because it is unneeded and can corrupt the initial ramdisk file system.

On SUSE, you might need to create an even larger initial ramdisk (for example, 48 MB or $48 \times 1024 \times 1024$ would result in a `seek=50331648`). If the initrd file is sufficiently large, skip ahead to mount the loopback device. See step 9.

6. Set up a loopback device for the image by typing the following command:

```
losetup /dev/loop0 /path/to/your/_initrd_file_
```

7. Ensure that you have a clean file system by again typing the following command:

```
e2fsck -f /dev/loop0
```

If you still have errors on the file system, the previous dd step was not performed correctly and it corrupted the initrd. You now must delete the loopback device by typing `losetup -d /dev/loop0` and restart the procedure from the beginning.

8. Resize the file system by typing the following command:

```
resize2fs /dev/loop0
```

Note: Resizing automatically expands the file system so that it uses all available space.

9. Mount the loopback device by typing the following command:

```
mount /dev/loop0 /mnt/tmp
```

10. You now have an initrd file system that is 32 MB. You can add additional files by typing the following command:

```
cd /mnt/tmp
```

11. If you have not already added your host adapter driver to the initrd file using the standard mk_initrd or mkinitrd process (depending on your distribution), you must manually copy the module files for the host adapter driver. You also must manually copy the SCSI core and SCSI disk drivers into the initrd file system and add the appropriate insmod command to the linuxrc script.
12. On SUSE, you must create the etc/, proc/, and sysroot/ directories in the initrd file system.
You can also add echo commands into the linuxrc script after the host adapter load and mounting /proc to force the addition of LUNs through /proc/scsi/scsi if the device discovery is not occurring automatically.
13. Create SDD directories in the initrd file system by typing the following commands:

```
mkdir -p opt/IBMsdd/bin  
chmod -R 640 opt/IBMsdd/
```

14. For SDD, you must copy the following files to the initrd file system.

Note: Ensure that you are in the /mnt/tmp directory when you perform the copies.

File Names	Target Location
/etc/vpath.conf	etc/
/etc/group	etc/
/etc/passwd	etc/
/etc/nsswitch.conf	etc/
/opt/IBMsdd/sdd-mod.o-CORRECT_VERSION	lib/sdd-mod.o
/opt/IBMsdd/bin/*	opt/IBMsdd/bin/
/lib/libc.so.6	lib/
/lib/ld-linux.so.2	lib/
/lib/libacl.so.1	lib/
/lib/libattr.so.1	lib/
/lib/libdl.so.2	lib/
/lib/libm.so.6	lib/
/lib/libpthread.so.0	lib/
/lib/libnss_files.so.2	lib/
/lib/librt.so.1	lib/
/bin/awk, chmod, chown, cp, date, grep, ls, mknod, mount, ps, rm, sed, sh, tar, unmount	bin/

File Names	Target Location
/dev/sd[a-z], sd[a-z][a-z] For example, tar cps /dev/sd[a-z] /dev/sd[a-z][a-z] tar xps	dev/

15. For Red Hat, you must copy the following additional files to the file system:

File Names	Target Location
/lib/libproc.so.2.0.7	lib/
/lib/libpcre.so.0	lib/
/lib/libtermcap.so.2	lib/
/bin/ash.static	bin/ash

16. For SUSE, you must copy the following additional files to the file system:

File Names	Target Location
/lib/libreadline.so.4	lib/
/lib/libhistory.so.4	lib/
/lib/libncurses.so.5	lib/
etc/nsswitch.conf Note: The etc/nsswitch.conf file must have its password and group entries changed to point to files instead of compat.	N/A

17. The following changes must be make to the initrd linuxrc script:

- a. For Red Hat, remove the following block of commands from the end of the file:

```
echo Creating block devices
mkdevices /dev
echo Creating root device
mkroot dev /dev/root
echo 0x0100 > /proc/sys/kernel/real-root-dev
echo Mounting root filesystem
mount -o defaults -ro - -t ext2 /dev/root /sysroot
pivot_root /sysroot /sysroot/initrd
unmount /initrd/proc
```

You must change the first line of the linuxrc script to invoke the ash shell instead of the nash shell.

- b. If the /proc file system is not already explicitly mounted in the linuxrc script, append the following mount command:

```
mount -n -tproc /proc /proc
```

- c. To configure SDD, append the following commands to the end of the linuxrc script:

```
insmod /lib/sdd-mod.o
/opt/IBMsdd/bin/cfgvpath
```

Mount the systems root file system so that you can copy configuration information to it. For example, if you have an ext3 root file system on

/dev/vpatha3, enter /bin/mount -o rw -t ext3, or /dev/vpatha3 /sysroot, or for a reiserfs root file system on /dev/vpatha3, enter /bin/mount -o rw -t reiserfs /dev/vpatha3 /sysroot

To copy the dynamically created device special files onto the systems root file system, enter the following commands:

```
tar cps /dev/IBMsdd /dev/vpath* | (cd /sysroot && tar xps)
/bin/umount /sysroot
```

You must define the root file system to the kernel. Traditionally, this information is passed to the bootloader as a string, for example /dev/vpatha3 and translated to a hexadecimal representation of the major and the minor numbers of the device. If the major and minor numbers equals 254,3, these numbers are represented in hex as 0xFE03. The linuxrc script passes the hexadecimal value into /proc with the following commands:

```
echo 0xFE03 > /proc/sys/kernel/real-root-dev
/bin/umount /proc
```

18. Edit the system fstab and change all the system mount points from LABEL or /dev/sd mount sources to their equivalent /dev/vpath. Refer to step 23 for dangers of booting by label in a multipath configuration.
19. Copy the system fstab to the initrd etc/ directory.
20. Unmount the image and remove the loopback binding by typing the following commands:

```
umount /mnt/tmp
losetup -d /dev/loop0
```

21. Compress the image by typing the following commands:

```
gzip -9 _initrd_file_
mv _initrd_file_.gz _initrd_file_
```

22. Append the following to your boot parameters (for example, in lilo.conf, grub.conf, or menu.lst):

```
ramdisk_size=34000
```

If you created a larger initrd file system, make this value large enough to cover the size.

23. For completeness, change the bootloader append for the root parameter of the kernel to the appropriate vpath device. For example, root=/dev/vpatha5. However, the previous steps override this value by passing the corresponding hex major, minor into the /proc file system within the initrd linuxrc script.

Note: If you boot by LABEL, there is a risk that the first device that is found in the fabric with the correct label could be the wrong device or that it is an sd single-pathed device instead of a vpath multipathed device.

24. Reboot the server. It will boot with the root file system on a vpath device instead of on an sd device.

Figure 72 on page 137 and Figure 73 on page 138 illustrate a complete linuxrc file for Red Hat and for SUSE.


```
#!/bin/ash

echo "Loading scsi_mod.o module"
insmod /lib/scsi_mod.o
echo "Loading sd_mod.o module"
insmod /lib/sd_mod.o
echo "Loading qla2310.o module"
insmod /lib/qla2310.o
echo "Loading jbd.o module"
insmod /lib/jbd.o
echo "Loading ext3.o module"
insmod /lib/ext3.o
echo Mounting /proc filesystem
mount -t proc /proc /proc
insmod /lib/sdd-mod.o
/opt/IBMsdd/bin/cfgvpath
/bin/mount -o rw -t ext3 /dev/vpatha3 /sysroot
tar cps /dev/IBMsdd /dev/vpath* | (cd /sysroot && tar xps)
/bin/umount /sysroot
echo 0xFE03 > /proc/sys/kernel/real-root-dev
/bin/umount /proc
```

Figure 72. Example of a complete linuxrc file for Red Hat

```

#!/bin/ash

export PATH=/sbin:/bin:/usr/bin

# check for SCSI parameters in /proc/cmdline
mount -n -tproc none /proc
for p in `cat /proc/cmdline` ; do
    case $p in
        scsi*|*_scsi_*|llun_blkfst=|max_report_luns=)
            extra_scsi_params="$extra_scsi_params $p"
            ;;
    esac
done
umount -n /proc

echo "Loading kernel/drivers/scsi/scsi_mod.o $extra_scsi_params"
insmod /lib/modules/2.4.21-190-smp/kernel/drivers/scsi/scsi_mod.o $extra_scsi_params

echo "Loading kernel/drivers/scsi/sd_mod.o"
insmod /lib/modules/2.4.21-190-smp/kernel/drivers/scsi/sd_mod.o

echo "Loading kernel/drivers/scsi/qla2310_conf.o"
insmod /lib/modules/2.4.21-190-smp/kernel/drivers/scsi/qla2310_conf.o

echo "Loading kernel/drivers/scsi/qla2310.o"
insmod /lib/modules/2.4.21-190-smp/kernel/drivers/scsi/qla2310.o

echo "Loading kernel/drivers/scsi/aic7xxx/aic7xxx.o"
insmod /lib/modules/2.4.21-190-smp/kernel/drivers/scsi/aic7xxx/aic7xxx.o

echo "Loading kernel/fs/reiserfs/reiserfs.o"
insmod /lib/modules/2.4.21-190-smp/kernel/fs/reiserfs/reiserfs.o
mount -t proc /proc /proc
insmod /lib/sdd-mod.o
/opt/IBMsdd/bin/cfgvpath
/bin/mount -o rw -t reiserfs /dev/vpatha3 /sysroot
tar cps /dev/IBMsdd /dev/vpath* | (cd /sysroot && tar xps)
/bin/umount /sysroot
echo 0xFE03 > /proc/sys/kernel/real-root-dev
/bin/umount /proc

```

Figure 73. Example of a complete linuxrc file for SUSE

Chapter 12. Attaching to an Intel host running VMware ESX server

This chapter provides instructions, requirements, and considerations for attaching an Intel host system running VMware ESX server to a storage unit.

Supported adapter cards for an Intel host running VMware ESX Server

This section provides a list of the supported adapter cards for an Intel host running VMware ESX Server.

This section describes how to attach an Intel server running VMware to a storage unit with the following fibre-channel adapters:

- QLogic QLA2310FL adapter card
- QLogic QLA2340FL adapter card
- QLogic QLA2342FL adapter card
- Emulex LP9402DC adapter card
- Emulex LP9802 adapter card

For up-to-date information about the servers, operating systems, distribution levels, host adapters, and fabric connections that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Note: The steps to install and configure adapter cards are examples. Your configuration might be different.

Attaching with fibre-channel adapters

This section describes how to attach an Intel server that is running a VMware ESX 2.5.1 server or a later server to an IBM storage unit.

For information about the most current version of the switches that are supported, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Attachment requirements for an Intel host running VMware ESX server

This section provides attachment requirements for an Intel host running VMware ESX Server.

- Check the LUN limitations according to the following information.

Number of paths (either more HBAs on the host or more ports on storage)	Maximum number of LUNs configured on storage / HBA
1	64
2	64
4	64
8	64
16	64
32	31

Note: The limitation calculated is a combination of the VMware limitations:

- Maximum of 64 LUNs running I/Os concurrently
- Less than 1024 LUNs managed by multipathing
- For details about the release level for your operating system, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Either you or an IBM service support representative must perform the following tasks to install and configure a storage unit:

1. Install the storage unit.
2. Define the fibre-channel host system with the worldwide port name identifiers. For the list of worldwide port names, see “Locating the worldwide port name (WWPN),” on page 225.
3. Define the fibre-port configuration if you did not do it when you installed the storage unit or fibre-channel adapters.
4. Configure the host system for the storage unit by using the instructions in your host system publications.

Attachment considerations for an Intel host running VMware ESX Server

This section provides attachment considerations for an Intel host running VMware ESX Server.

For the most current VMware ESX versions, servers, and switches, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

See <http://www.vmware.com> to download the vendor updates.

Installing the Emulex adapter card for an Intel host running VMware ESX Server

This section provides instructions for installing the Emulex adapter drivers for an Intel host running VMware ESX Server.

Note: The LP9402DC adapter card is a 3.3 volt card only and requires a 3.3 volt PCI slot.

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

- Target
 - Public initiator
 - Private initiator
 - Target and public initiator
 - Target and private initiator
1. Record the Institute of Electrical and Electronics Engineers (IEEE) number that is printed on the card. You can use the IEEE number to determine the WWPN.

Note: You can find the WWPN in your `/proc/scsi/<adapter_type>/<n>` directory, where `<adapter_type>` is the host adapter type and `<n>` the host adapter number for your card. Your host adapter type can be either `qlaxxxx` for QLogic adapters or `lpfccd` for Emulex adapters.

2. Refer to the installation instructions provided by your host adapter vendor for their specific installation instructions. See <http://www.emulex.com> for the latest documentation.

Installing the QLogic adapter card for an Intel host running VMware ESX Server

You can use these instructions to install the QLogic adapter card for an Intel host running VMware ESX Server.

This section tells you how to attach an Intel host system running VMware ESX Server to a storage unit with the QLogic QLA23xx adapter card. Single- and dual-port fibre-channel interfaces with the QLogic adapter cards support the following public and private loop modes:

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

Note:

1. The following steps are an example configuration. The configuration for your adapter might differ.
2. The complete list of host adapter settings for attachment to the storage unit can be found at:

<http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do>

1. Install the QLogic QLA23xx adapter card in the host system.
2. Connect the cable to the storage unit port.
3. Restart the server.
4. Press Ctrl+Q to get to the FAST!Util menu panel.
5. Use FAST!Util to set the parameters to the values listed in the following table:

Table 9. Recommended settings for the QLogic adapter card for an Intel host running VMware ESX Server

Parameters	Recommended settings
Basic Input/Output System (BIOS)	Disabled
Frame size	2048
Loop reset delay	5 (minimum)
Adapter hard loop ID	Disabled
Hard loop ID	0
Execution throttle	100
Fast command posting	Enabled
>4 GB addressing	Disabled (for 32-bit systems)
LUNs per target	0 or 128
Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes

Table 9. Recommended settings for the QLogic adapter card for an Intel host running VMware ESX Server (continued)

Parameters	Recommended settings
Login retry count	<ul style="list-style-type: none"> • 60 if you are not running the SDD • 8 if you are running the SDD
Port down retry count	<ul style="list-style-type: none"> • 60 if you are not running the SDD • 8 if you are running the SDD
Driver load RISC code	Enabled
Enable database updates	No
Disable database load	No
IOCB allocation	256
Extended error logging	Disabled (You might want to enable this parameter if you use debugging.)
Type of connection	Point-to-point

6. Restart the server.

Defining the number of disks devices on VMware ESX Server

This section provides information for defining the number of disks devices on VMware ESX Server.

The maximum number of devices that are supported on a VMware ESX Server is 128. However the VMkernel supports up to 64 concurrent I/Os. The embedded multipathing capability of the VMkernel can manage up to 1024 disks.

The maximum number of LUNs that can be used must accommodate the 3 previous limitations.

SCSI disk considerations for an Intel host running VMware ESX server

This section provides information for identifying LUNs.

LUN identification for the VMware ESX console

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

Each of the attached storage unit LUNs has a special device file in the Console OS directory /dev. There is a maximum of 128 fibre-channel disks entries with the following format /dev/vsd [0 -127].

From the VMware ESX console, the storage unit can be traced through several layers:

- Host adapter level: the Host Adapter driver can provide the list of LUNs that are detected for each fibre-channel port.
- Vmhba level: VMware virtualizes the Host Adapter as a *vmhba* from which you can retrieve the information that is related to the multipathing.

Host adapter level

The Host Adapter driver (lpfc for Emulex) (qla23xx for QLogic) can provide information on the port base. The content of the information varies according to the

driver. The directory `/proc/scsi/ <driver name> / <adapter number>` will provide the information for a specific adapter/port. See Figure 74 for an example of QLogic output and Figure 75 on page 144 for an example of Emulex output.

```
[root@denethor root]# cat /proc/scsi/qla2310/2
QLogic PCI to Fibre Channel Host Adapter for ISP23xx:
    Firmware version: 3.01.18, Driver version 6.04.02
Entry address = 0x4d8000
HBA: QLA2312 , Serial# H88769
Request Queue = 0x210ec000, Response Queue = 0x21100000
Request Queue count= 128, Response Queue count= 512
Total number of active commands = 0
Total number of interrupts = 4008395
Total number of IOCBs (used/max) = (0/600)
Total number of queued commands = 0
    Device queue depth = 0x10
Number of free request entries = 128
Number of mailbox timeouts = 0
Number of ISP aborts = 90
Number of loop resyncs = 300
Number of retries for empty slots = 0
Number of reqs in pending_q= 0, retry_q= 0, done_q= 0, scsi_retry_q= 0
Host adapter:loop state= <DOWN>, flags= 0x860813
Dpc flags = 0x0
MBX flags = 0x0
SRB Free Count = 4096
Port down retry = 30
Login retry count = 30
Commands retried with dropped frame(s) =0

SCSI Device Information:
scsi-qla1-adapter-node=200000e08b0c2109;
scsi-qla1-adapter-port=210000e08b0c2109;
scsi-qla1-target-0=5005076300c6a3bb;
scsi-qla1-target-1=200500a0b80f0a15;
scsi-qla1-target-2=5005076300c6a6b8;
scsi-qla1-target-3=5005076300c2a6b8;
scsi-qla1-target-4=5005076300cfa6b8;

SCSI LUN Information:
(Id:Lun)
( 4: 0): Total reqs 7, Pending reqs 0, flags 0x0, 1:0:81,
( 4: 1): Total reqs 7132, Pending reqs 0, flags 0x0, 1:0:81,
( 4: 2): Total reqs 7414, Pending reqs 0, flags 0x0, 1:0:81,
( 4: 3): Total reqs 2948, Pending reqs 0, flags 0x0, 1:0:81,
...
```

Figure 74. Example of QLogic Output:

```

HBA: Emulex LightPulse LP9002 on PCI bus 06 device 28 irq 49
Emulex LightPulse LPFC Driver Version: 4.20q
SerialNum: MI23710129
Firmware Version: 3.90A7
Hdw: 2002606d
VendorId: 0xf90010df
Portname: 10:00:00:00:c9:2e:78:cc  Nodename: 20:00:00:00:c9:2e:78:cc

Link Up - Ready:
  PortID 0x71000a
  Fabric
  Current speed 2G

lpfc1t00 DID 71000e WWPN 50:05:07:63:00:cf:a6:b8 WWNN 50:05:07:63:00:c0:a6:b8

```

Figure 75. Example of Emulex Output:

Vmhba level

The `/proc/vmware/scsi` directory will present 1 directory `vmhba<N>` where `<N>` is the host adapter. The directory will contain 1 entry per device (with a format `<scsi id> : <LUN number>`). See Figure 76 for an example of a Vmhba directory.

```

[root@denethor root]# ls /proc/vmware/scsi/vmhba0/
4:0  4:14 4:2  4:25 4:30 4:36 4:41 4:47 4:52 4:58 4:63 4:69 stats
4:1  4:15 4:20 4:26 4:31 4:37 4:42 4:48 4:53 4:59 4:64 4:7
4:10 4:16 4:21 4:27 4:32 4:38 4:43 4:49 4:54 4:6  4:65 4:70
4:11 4:17 4:22 4:28 4:33 4:39 4:44 4:5  4:55 4:60 4:66 4:71
4:12 4:18 4:23 4:29 4:34 4:4  4:45 4:50 4:56 4:61 4:67 4:8
4:13 4:19 4:24 4:3  4:35 4:40 4:46 4:51 4:57 4:62 4:68 4:9

```

Figure 76. Example listing of a Vmhba directory

Note: In a multipathing configuration, only 1 vmhba directory displays entries for the shared device.

Each of the entries provides information on the drives that it points to, as shown in Figure 77 on page 145:


```

[root@denethor root]# cat /proc/vmware/scsi/vmhba0//4\:71
Vendor: IBM          Model: 2105800      Rev: 1.89
Type:  Direct-Access          ANSI SCSI revision: 03
Size:  24795 Mbytes
Queue Depth: 16

Partition Info:
Block size: 512
Num Blocks: 50781312

      num:      Start      Size      Type
      1:         63 50765337      fb

Partition      VM      cmds      reads      KBread      writes      KBwritten      cmdsAbt
busRst      paeCmds      paeCopies      splitCmds      splitCopies      issueAvg      totalAvg
0          -          168          87          46          0          0 0          0
0          0          0          0          0          20788          936149
1          -          8837          8581          33869          128          64 0          0
0          0          0          0          0          14038          762180

      VM      Shares      cmds      reads      KBread      writes      KBwritten      cmdsAbt      busRst
paeCmds      paeCopies      splitCmds      splitCopies      issueAvg      totalAvg      active      queued
virtTime
127      1000          9005          8668          33915          128          64          0          0
0          0          0          0          0          14164          765426          0          0
8728000000
Total      1000          9005          8668          33915          128          64          0          0
0          0          0          0          0          14164          765426          0          0
8728000000

Paths:fixed
vmhba0:4:71 on*#
vmhba1:4:71 on

Active: 0  Queued: 0

```

Figure 77. Example of Vmhba entries

Note: The last section displays information on multipathing; refer to VMware documentation for more information on the multipathing configuration. The storage unit can interact with both FIXED and MRU policies.

Disk device discovery on VMware ESX

This section provides information about disk device discovery on VMware ESX.

To rediscover LUNs from the Console OS use the command `vmkfstools -s vmhbaC` where C is the number of the virtual host adapter to rescan.

Persistent binding

This section provides information about persistent binding.

Persistent binding is highly recommended in a SAN environment especially when several storage subsystems are available to the VMware ESX Server. The storage SCSI ID, in a SAN, is dictated by the order in which the VMware ESX Server will discover the subsystems at boot time. In some instances, this order might defer because of a device addition or removal. This could cause VMware ESX Server to assign a different Target ID and change the device addressing (vmhba <C> : <T> : <D>). This can cause the virtual machines to report their storage as missing when booting. Using the persistent binding will force VMware ESX Server to always use the same Target ID for a storage unit no matter of the discovery order.

To enable persistent binding on all adapter use the command: `pbind.pl -A`

Configuring the storage unit

This section provides instructions for configuring the storage unit.

Note: IBM does not support VMware ESX running on “system LUN/disk” or “extended VMFS2” file system.

Configuring the storage unit for an Intel host running VMware ESX Server

This section provides instructions for configuring the storage unit for an Intel host running VMware ESX Server.

VMware ESX multipathing requires that the storage unit present the disk devices in the same LUN sequence for each of the host's fibre-channel adapters. Therefore, when “assigning volumes to target hosts” using IBM TotalStorage DS Storage Manager to the host adapters, it is necessary to select the option “use same ID/LUN in source and target”.

Note: This requirement is only necessary when the different host adapters are located on the same VMware ESX Server, not when devices are shared between VMware ESX Servers.

Partitioning storage unit disks for an Intel host running VMware ESX Server

This section provides information about partitioning storage unit disks for an Intel host running VMware ESX Server.

Before you create a VMFS file system, partition the disk by using the `fdisk` utility. You have to specify the special device file of the disk that you want to partition when you run `fdisk`. The type of the partition that will be used by VMFS has to be `fb` (must be set with `t` command of `fdisk`).

To identify the special device file to use, you must run the command: `vmkfstools -N vmhba<C>:<T>:<D>:0` where `<C>` is the virtual host adapter number, `<T>` is the SCSI ID of the storage unit, and `<D>` is the LUN number of the disk.

```
[root@denethor root]# vmkfstools -N vmhba0:4:60:0  
/dev/vsd60
```

```
[root@denethor root]# fdisk /dev/vsd71
```

The number of cylinders for this disk is set to 3160.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:

- 1) software that runs at boot time (e.g., old versions of LILO)
- 2) booting and partitioning software from other OSs
(e.g., DOS FDISK, OS/2 FDISK)

Command (m for help): m
Command action

- a toggle a bootable flag
- b edit bsd disklabel
- c toggle the dos compatibility flag
- d delete a partition
- l list known partition types
- m print this menu
- n add a new partition
- o create a new empty DOS partition table
- p print the partition table
- q quit without saving changes
- s create a new empty Sun disklabel
- t change a partition's system id
- u change display/entry units
- v verify the partition table
- w write table to disk and exit
- x extra functionality (experts only)

Figure 78. Example of the different options for the fdisk utility:

```
Command (m for help): n
Command action
  e extended
  p primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-3160, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-3160, default 3160):
Using default value 3160

Command (m for help): p

Disk /dev/vsd71: 255 heads, 63 sectors, 3160 cylinders
Units = cylinders of 16065 * 512 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vsd71p1		1	3160	25382668+	83	Linux

```

Command (m for help): t
Partition number (1-4): 1
Hex code (type L to list codes): fb
Changed system type of partition 1 to fb (Unknown)

Command (m for help): p

Disk /dev/vsd71: 255 heads, 63 sectors, 3160 cylinders
Units = cylinders of 16065 * 512 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vsd71p1		1	3160	25382668+	fb	Unknown

Figure 79. Example of primary partition on the disk /dev/vsd71

Creating and using VMFS on the storage unit for an Intel host running VMware ESX Server

This section provides information for creating and using VMFS on the storage unit for an Intel host running VMware ESX Server.

After you partition the disk as described in “Partitioning storage unit disks for an Intel host running VMware ESX Server” on page 146, the next step requires that you create a VMFS file system by using the following commands:

```
vmkfstools -C vmfs2 <device>
```

<device> is on the format vmhba <C>:<T>:<D>:<P> , where <C> is the virtual host adapter, <T> is the SCSI ID of the storage unit, <D> is the Lun ID, and <P> refers to the partition number to use on that device (previously created with fdisk). In most cases it will be 1.

```
[root@denethor vmfs]# vmkfstools -C vmfs2 vmhba0:4:71:1
This is already a file system on vmhba0:4:71:1. Overwrite it (y/n)? y
Creating file system on vmhba0:4:71:1 with blockSize 1048576
```

After this operation is completed, the new file system is mounted under /vmfs.

```
vmkfstools -N <label name>
```

Note: This step is optional but we strongly recommended that you do it because it can ease your task of storage management. It can become a mandatory step if you do not configure persistent binding on VMware ESX Server.

Copy Services considerations

This section provides considerations for using copy services.

Copy Services are only available for the VMware ESX Console OS

IBM supports FlashCopy and Remote Mirror and Copy between different VMware ESX Servers with the following limitations:

- Copy Services are only available for the VMware ESX Console OS.
- The source and target volumes must exist on separate physical machines.
- Remote Mirror and Copy read from secondary option is not supported.
- Rescan of LUNs might be required.

IBM does not support “system LUN/disk” for the virtual machine. Therefore, you must use VMFS 1 or 2. The locking mechanism of VMFS might get confused when performing copy services on the target volume. To prevent this, it is mandatory that no virtual machine using a copy target volume is running during the copy process.

The source and target volumes must be on separate physical machines

VMFS file system uses a label mechanism and if it encounters 2 volumes with the same label, only one remains accessible.

Remote Mirror and Copy read from secondary option is not supported

“Read from secondary” allows a target volume to be opened in read-only mode during a Remote Mirror and Copy operation. VMware ESX Server does not open volumes in read-only mode and therefore is not able to access a Remote Mirror and Copy target while a copy is in progress.

Rescan of LUNS might be required

VMFS is able to manage (without any “umount”) changes on a hard drive at the file system level. It can manage a copy as long as the system knows that there is a VMFS file system to look for:

- If the copy target volume was discovered at boot time without a VMFS file system on it, you must perform a rescan operation for the changes.
- If it knows that there is a VMFS file system, the server can use the new information immediately without any intervention.

Chapter 13. Attaching to a Novell NetWare host

This chapter provides instructions, requirements, and considerations for attaching a Novell NetWare host system to a storage unit.

Attaching with fibre-channel adapters

This section describes how to attach a storage unit to a Novell NetWare host system with fibre-channel adapter cards.

Attaching a Novell NetWare host with fibre-channel adapters

This section lists the fibre-channel adapters for Novell NetWare hosts.

This section describes how to attach a storage unit to a Novell NetWare host system with the following adapter cards:

- Emulex LP9002L adapter card
- Emulex LP9002DC adapter card
- Emulex LP9402DC adapter card
- Emulex LP9802 adapter card
- Emulex LP 10000 adapter card
- Emulex LP 10000DC adapter card
- QLogic QLA2310F adapter card
- QLogic QLA2310FL adapter card
- QLogic QLA2340FL adapter card
- QLogic QLA2342FL adapter card

For a list of open-systems hosts, operating systems, adapters and switches that IBM does not support, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Fibre-channel attachment considerations for a Novell NetWare host

Fibre-channel attachment considerations for a Novell NetWare host

For up to date information about the switch and fabric connections that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Installing the Emulex adapter card for a Novell NetWare host

This section tells you how to attach a storage unit to a Novell NetWare host system with an Emulex adapter card.

Note: The Emulex LP9402DC adapter card is a 3.3 volt card only and requires a 3.3 volt PCI slot.

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

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

- Target and private initiator
- 1. Record the IEEE number that is printed on the card. You can use the IEEE number to determine the WWPNN.
- 2. Refer to the installation instructions provided by your host adapter vendor for their specific installation instructions. See <http://www.emulex.com> for the latest documentation.

Downloading the current Emulex adapter driver for a Novell NetWare host

This section provides instructions for downloading the current Emulex adapter driver for a Novell NetWare host.

1. Plug in and restart your host system.
2. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
3. Click **Interoperability matrix**.
4. Click **DS8000 interoperability matrix**.
5. Find the section for the current version of the driver and firmware and driver you want.
6. Go to <http://www.emulex.com>
7. Click **Drivers, software and manuals** from the left navigation pane
8. Select **IBM** from the drop down vendor list.
9. Click the link for the name and version of the adapter driver you want to download.

Installing the Emulex adapter driver for a Novell NetWare host

This topic provides installation steps for installing the Emulex adapter driver on a Novell NetWare host.

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

1. Go to www.emulex.com.
2. Click **drivers, software and manuals** in the left navigation pane.
3. From the **Drivers, software and manuals by operating system** section, click **NetWare**.
4. Click **Installation** to get the documentation that tells you how to install the adapter driver.
5. Once you have installed the driver from the NetWare server console, enter the following command: `nwconfig`. This command applies to NetWare 5.0 or later.
6. Click **NCF Files Options**.
7. Select **Edit STARTUP.NCF** file.
8. Insert the following statement at the beginning of the STARTUP.NCF file: `SET MULTI-PATH SUPPORT=OFF`
9. Find the LOAD statement for the LPFC.HAM driver and ensure that it reads as follows: `LOAD LPFC.HAM SLOT=XXXX TOPOLOGY=1 /LUN MAX_LUN=32 RSCNZONE=1`

Downloading the current QLogic adapter driver for a Novell NetWare host

This section provides instructions for downloading the current QLogic adapter driver for a Novell NetWare host.

1. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

2. Click **Interoperability matrix**.
3. Click **DS8000 interoperability matrix**.
4. Find the section for the current version of the driver and firmware and driver you want.
5. Go to <http://www.qlogic.com>
6. Click **Drivers / Software** from the **Quick Links** list.
7. Choose **IBM** from the **OEM and ISV approved Drivers and Firmware** section.
8. Navigate to the adapter you want to download, then click **Download**.

Installing the QLogic QLA23xx adapter card for a Novell NetWare host

This section tells you how to attach a storage unit to a Novell NetWare host system with the QLogic QLA23xx adapter card. Single- and dual-port fibre-channel interfaces with the QLogic QLA23xx adapter card support the following public and private loop modes.

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

Note: The following steps are an example of a configuration. The configuration for your adapter might differ.

1. Install the QLogic QLA23xx adapter card in the host system.
2. Connect the cable to the storage unit port.
3. Restart the server.
4. Press Ctrl+Q to get to the **FAST!Util** menu.
5. Use FAST!Util to set the parameters to the values listed in the following table:

Table 10. Recommended settings for the QLogic QLA23xx adapter card for a Novell NetWare host

Parameters	Recommended settings
Basic Input/Output System (BIOS)	Disabled
Frame size	2048
Loop reset delay	5 (minimum)
Adapter hard loop ID	Disabled
Hard loop ID	0
Execution throttle	16
Fast command posting	Enabled
>4 GB addressing	Disabled (for 32-bit systems)
LUNs per target	32 Note: Set LUNs per target to a number greater than the LUNs you have assigned to the adapter.
Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes
Login retry count	20 (minimum)

Table 10. Recommended settings for the QLogic QLA23xx adapter card for a Novell NetWare host (continued)

Parameters	Recommended settings
Port down retry count	20 (minimum)
Driver load RISC code	Enabled
Enable database updates	No
Disable database load	No
IOCB allocation	256
Extended error logging	Disabled (You might want to enable this parameter if you use debugging.)
Type of connection	<ul style="list-style-type: none"> • Loop only • Point-to-point • Loop preferred (If you cannot use arbitrated loop, then default to point-to-point.) • Point-to point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop.) <p>Note: If you connect the storage unit directly to the host system, the option you select must match the port connections on the storage unit.</p>

6. Restart the server.

Installing the adapter drivers for a Novell NetWare host

This section provides instructions for installing the adapter drivers for a Novell NetWare host.

1. From the NetWare server console, type `nwconfig` for NetWare 5.0 or later.
2. Click **Driver Options**.
3. Click **Configure Disk and Storage Device Drivers**.
4. Click **Select an additional driver**.
5. Press the Insert key.
6. Insert a floppy diskette with the QLogic drivers into the A:\ drive of the NetWare server. Press Enter.
The available driver is displayed.
7. Click the driver for the QLogic card, and press Enter.
8. Select **Select/Modify driver parameters**, and type the slot number of the QLogic card into the slot number parameter.
9. Set the Scan All LUNs parameter to Yes.
10. Press the Tab key and check the **Save Parameters and Load Driver** check box.
11. Press Esc twice to return to the **Configuration Options** menu.
 - a. Highlight the **NCF Files Options** menu and press Enter.
 - b. Select **Edit the STARTUP.NCF** and press Enter.
 - c. Verify the location of the STARTUP.NCF file, and modify it as necessary. Press Enter.

- d. Enter the following statement at the beginning of the STARTUP.NCF file:
SET MULTI-PATH SUPPORT=OFF
 - e. Find the LOAD statement for the QL2300.HAM driver and ensure that it reads as follows:
LOAD QL2300.HAM slot = xxxx/luns/xretry = 12/xtimeout = 120
 - f. Find the line that reads:
LOAD QLCSFTE
REM LOAD QLCSFTE
 - g. Press Esc and select **Yes** when prompted to save the file.
12. Type UNLOAD QL2300.HAM.
13. Type LOAD QL2300.HAM slot = xxxx/luns/xretry = 12/xtimeout = 120, where xxxx is the PCI slot number where the QL23xx card is installed.

The storage unit volumes are displayed in the devices list. Create volumes using the nssmu utility if necessary.

Chapter 14. Attaching to a iSCSI Gateway host

This chapter describes how to attach a storage unit to a iSCSI Router host system with the following adapter cards:

- For Ethernet using the TCP/IP protocol:
 - Netfinity® 1000SX 1GbE adapter (fibre-channel connection)
 - Alacritech 100/1000 1GbE adapter (CAT5 RJ45 connection)
 - Integrated 10/100 Ethernet connections (CAT5 RJ45 connection)
 - NetGear GA600

For up to date information about the servers, operating systems, host adapters, and fabric connections that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Note: The steps to install and configure adapter cards are examples. Your configuration might be different.

Attachment considerations

This section provides an overview for attaching a storage unit to a iSCSI Gateway host.

Attachment overview of the iSCSI Gateway host

This topic provides an attachment overview of the iSCSI Gateway host.

The iSCSI Gateway router provides access to storage unit fibre-channel storage over an Ethernet (LAN based) network. Host servers attached through Ethernet connections can attach to assigned LUNs in much the same way as fibre-channel attached servers by using host bus adapters. LUNs can shared or dedicated.

Note: When you configure the storage unit for host attachment, you can define the worldwide port name by host type. For example, you can define the SN5420 as a iSCSI Gateway host type. You can identify each host by an entry in the scsirouter list. You can present the storage unit LUNs to the SN5420, and you can use the scsirouter list to manage a host assignment.

Attachment requirements for the iSCSI Gateway host

This section provides attachment requirements for the iSCSI Gateway host.

This section lists the requirements for attaching the storage unit to your host system:

- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *User's Guide* is on the TotalStorage CD that you receive with the storage unit.
- See <http://www-1.ibm.com/servers/storage/disk/ds8000> for details about the release level for your operating system.

Either you or an IBM service support representative must perform the following tasks to install and configure a storage unit:

1. Install the storage unit.

2. Define the fibre-channel host system with the worldwide port name identifiers. For the list of worldwide port names, see “Locating the worldwide port name (WWPN),” on page 225.
3. Define the fibre-port configuration if you did not do it when you installed the storage unit or fibre-channel adapters.
4. Configure the host system for the storage unit by using the instructions in your host system publications.

A single SN5420 provides a 100 Mbps connection to a storage area network. You can define up to four SCSI routers for the SN5420. You can have up to thirty-two hosts on each SCSI router. Each host can have up to eight storage targets. Each target can present multiple LUNS. For example, you can define eight LUNS for Windows 2000 or Windows 2003. You can define thirty-two LUNS for the UNIX operating system or the Linux operating system. A Windows iSCSI host can discover sixty-four storage devices or eight targets times eight LUNs. A UNIX or Linux host can discover 256 Storage devices or eight targets times thirty-two LUNs.

Ethernet adapter attachment considerations for the iSCSI Gateway host

This topic provides Ethernet adapter attachment considerations for the iSCSI Gateway host.

The iSCSI software driver resides on each host (initiator) and completes the connection to a storage device (target). You can only use the iSCSI software drivers (initiators) for the following host systems:

- Windows 2000
- Windows 2003
- Linux (Red Hat Enterprise Linux Server 2.1 and above)

The iSCSI protocol is based on TCP/IP. The iSCSI protocol depends on TCP/IP addressing to identify the SCSI router lists and storage targets. You must install an Ethernet adapter Network Interface Card (NIC) in the host computer.

You can connect NICs of 10/100 or 10/100/1000 through CAT5 copper or fibre connector cables.

Note: For a list of open-systems hosts, the operating systems, adapters and switches that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Configuring for storage for the iSCSI Gateway host

This topic provides information for configuring storage for the iSCSI Gateway host.

You must configure the storage unit to present storage LUNs to the storage router. You must configure the storage router to list the storage LUNs in the SCSI router definition. You must assign a TCP/IP address to the SCSI router. You must identify an entry in the access list table for the storage router for each host. You can group hosts that are listed in the access list table by operating system. The iSCSI software driver must have the TCP/IP address that points to the scsirouter target, and the list of storage LUNs defined in the scsirouter entry.

iSCSI Gateway operation through the IP Service Module

This topic discusses the support provided by the IP Service Module.

iSCSI Gateway support is available through the use of the IP Service Module. The IP Service Module supports the following host systems:

- Intel hosts running the Red Hat Enterprise Linux Advanced Server 2.1 and above operating system
- Intel hosts running the Microsoft Windows 2000 operating system
- Intel hosts running the Microsoft Windows Server 2003 operating system

Chapter 15. Attaching to an IBM SAN File System

This chapter provides the information that you need to attach a storage unit to IBM SAN File System metadata servers.

Attaching to an IBM SAN File System metadata server with Fibre-channel adapters

This section provides information about attaching an IBM SAN File System with fibre-channel adapters.

You can attach a storage unit to the IBM SAN File System metadata engine using the SAN File System fibre-channel adapter. You cannot attach a storage unit to an IBM SAN File System host system using SCSI adapters

Note: For a list of open-systems hosts, operating systems, adapters and switches that IBM supports, see <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>. For additional information about the IBM SAN File System, see <http://www-1.ibm.com/servers/storage/support/virtual/sanfs.html>

Configuring a storage unit for attachment to the SAN File System metadata node

This section provides instructions for configuring a storage unit for attachment to the SAN File System metadata engine.

You must configure the storage unit before you can successfully attach it to SAN File System metadata engines.

Perform the following steps to configure the storage unit for attachment to the SAN File System metadata engine. You must create a host entry for each port on every SAN File System metadata engine in the cluster. For example, a four-node SAN File System cluster requires 8 host entries.

1. Start the DS Storage Manager on the management console.
2. In the navigation, under Manage Hardware, select Host Systems. In the Host Systems – Main Page, select the desired storage complex and storage unit and then select **Create...** in the **Select Action** drop-down list. Click **Go** to display the Create Host System – General host information page.
3. Complete the fields in the General host information page.
 - a. Select **IBM SAN File System – Linux MDS(SanFsLinux)** as the host type.
 - b. Enter the host nickname.
 - c. You can optionally enter a host description.
 - d. Click **Next**. The Create Host System – Define host ports page is displayed.
4. Complete the fields in the Define host ports page.
 - a. Specify the quantity.
 - b. Specify the attachment port type.
 - c. Add at least one host port definition to the Defined host ports table.

- d. You can optionally check the **Group ports to share a common set of volumes** box. This causes the quantity of ports that are identified in the Quantity field to become grouped together and treated as a single host attachment.
 - e. Select at least one host port definition from the Defined host ports table.
 - f. Click **Next**. The Create Host System – Specify Host Port WWPN page is displayed.
5. In the Specify Host Port WWPN page, use one of the following two methods to specify the host port WWPNs:
 - Select the WWPN from the list. To do this, you must have already zoned the SAN File System with a storage unit, and the SAN File System must be turned on.
 - Manually enter the WWPN for the host port.

Note: A configuration command fails if you use a WWPN that equals zero (0).

Click **Next**. The Create Host System – Specify storage units page is displayed.

6. Specify the storage units for the host attachment field by selecting a storage unit from the Available storage units list and clicking **Add**. Click **Next** to display the Create Host System – Specify storage unit parameters page.

Note: If you select the **Create a Storage Unit** button (Simulated only), follow the process for creating the new storage unit. Once you have completed that process by clicking the **Finish** button, the new storage unit is available for selection.

7. Complete the fields in the Specify storage unit parameters page.
 - a. Select a host attachment ID.
 - b. Select a storage group.
 - c. Choose a login option.

You can loop through this page for each host attachment identifier by clicking the **Apply Assignment** button to commit the current transaction and then starting the procedure again by selecting another identifier. If you select an existing host attachment identifier from the table, you can click the **Create a new group** button to create a new volume group for selection. If you decide that this host attachment can log in to **The following specific storage image ports**, you must specify the specific ports in the Available I/O storage ports table.

When you are finished specifying the storage unit parameters, click **Ok**. The Create Host System – Verification page is displayed.

8. In the Verification page, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and specify the correct values. Otherwise, click **Finish** to complete the host system creation process.

Chapter 16. Attaching to an IBM SAN Volume Controller host

This chapter provides the information that you need to attach a storage unit to an IBM SAN Volume Controller host system.

Attaching to an IBM SAN Volume Controller host with fibre-channel adapters

You can attach a storage unit to the IBM SAN Volume Controller host system using the SAN Volume Controller (SVC) fibre-channel adapter.

Note: For a list of open-systems hosts, operating systems, adapters and switches that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

For further information about the SAN Volume Controller host system, see <http://www-1.ibm.com/servers/storage/support/virtual/2145.html>.

Configuring the storage unit for attachment to the SAN Volume Controller host system

Perform the following steps to configure the storage unit for attachment to the SAN Volume Controller host system.

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

1. Start the DS Storage Manager on the management console.
2. In the navigation, under Manage Hardware, select Host Systems. In the Host Systems – Main Page, select the desired storage complex and storage unit and then select **Create...** in the **Select Action** drop-down list. Click **Go** to display the Create Host System – General host information page.
3. Complete the fields in the General host information page.
 - a. Select **IBM SAN Volume Controller (SVC)** as the host type.
 - b. Enter the host nickname.
 - c. You can optionally enter a host description.
 - d. Click **Next**. The Create Host System – Define host ports page is displayed.
4. Complete the fields in the Define host ports page.
 - a. Specify the quantity.
 - b. Specify the attachment port type.
 - c. Add at least one host port definition to the Defined host ports table.
 - d. You can optionally check the **Group ports to share a common set of volumes** box. This causes the quantity of ports that are identified in the Quantity field to become grouped together and treated as a single host attachment.
 - e. Select at least one host port definition from the Defined host ports table.
 - f. Click **Next**. The Create Host System – Specify Host Port WWPN page is displayed.
5. In the Specify Host Port WWPN page, use one of the following two methods to specify the host port WWPNs:

- Select the WWPN from the list. To do this, you must have already zoned the SAN Volume Controller with a storage unit, and the SAN Volume Controller must be turned on.
- Manually enter the WWPN for the host port.

Note: A configuration command fails if you use a WWPN that equals zero (0).

Click **Next**. The Create Host System – Specify storage units page is displayed.

6. Specify the storage units for the host attachment field by selecting a storage unit from the Available storage units list and clicking **Add**. Click **Next** to display the Create Host System – Specify storage unit parameters page.

Note: If you select the **Create a Storage Unit** button (Simulated only), follow the process for creating the new storage unit. Once you have completed that process by clicking the **Finish** button, the new storage unit is available for selection.

7. Complete the fields in the Specify storage unit parameters page.
 - a. Select a host attachment ID.
 - b. Select a storage group.
 - c. Choose a login option.

You can loop through this page for each host attachment identifier by clicking the **Apply Assignment** button to commit the current transaction and then starting the procedure again by selecting another identifier. If you select an existing host attachment identifier from the table, you can click the **Create a new group** button to create a new volume group for selection. If you decide that this host attachment can log in to **The following specific storage image ports**, you must specify the specific ports in the Available I/O storage ports table.

When you are finished specifying the storage unit parameters, click **Ok**. The Create Host System – Verification page is displayed.

8. In the Verification page, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and specify the correct values. Otherwise, click **Finish** to complete the host system creation process.

Chapter 17. Attaching an SGI host system

This chapter provides instructions, requirements, and considerations for attaching an SGI host system to a storage unit.

Attaching an SGI host system with fibre-channel adapters

This section lists the fibre-channel adapters supported for the SGI host system.

- PCI-FC-1P-OPT
- PCI-FC-1P-OPT-A
- PCI-FC-1P-OPT-B
- PCX-FC-2P-OPT-B

You can attach an SGI host system to a storage unit with fibre-channel adapters.

You can also connect a nonclustered SGI host system to the storage unit using the switched fabric topology on optical fiber media, SGI fibre-channel host bus adapters, and the IBM 2109 switch. You can also connect directly from an SGI host system to a storage unit with the arbitrated loop topology.

For current information about the servers, operating systems, host adapters, and fabric connections that IBM supports for an SGI host system, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Note: The steps to install and configure adapter cards are examples. Your configuration might be different.

Fibre-channel attachment considerations for the SGI host system

This section lists fibre-channel attachment considerations for the SGI host system.

Ensure that you have the IRIX64 6.5.14 version or above installed on your host system.

The SGI host system does not support IBM's Subsystem Device Driver (SDD) because IRIX manages multipathing.

Fibre-channel attachment requirements for the SGI host system

This section provides fibre-channel attachment requirements for the SGI host system.

- Check the LUN limitations for your host system.
- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *IBM TotalStorage DS8000 User's Guide* is on the TotalStorage CD that you receive with the storage unit.
- See <http://www.ibm.com/servers/storage/disk/ds8000/interop.html> for details about the release level for your operating system.

Either you or an IBM service support representative (SSR) must perform the following tasks to install and configure a storage unit.

1. The IBM SSR installs the storage unit by using the procedures in the *IBM TotalStorage DS8000 Service Guide*.
2. Either you or an IBM SSR defines the fibre-channel host system with the worldwide port name identifiers. For the list of worldwide port names, see “Locating the worldwide port name (WWPN),” on page 225.
3. Either you or an IBM SSR defines the fibre-port configuration if you did not do it during the installation of the storage unit or fibre-channel adapters.
4. Either you or an IBM SSR uses the instructions in your host system publications to configure the host system for the storage unit.

Either you or an IBM service support representative must perform the following tasks to install and configure a storage unit:

1. Install the storage unit.
2. Define the fibre-channel host system with the worldwide port name identifiers. For the list of worldwide port names, see “Locating the worldwide port name (WWPN),” on page 225.
3. Define the fibre-port configuration if you did not do it when you installed the storage unit or fibre-channel adapters.
4. Configure the host system for the storage unit by using the instructions in your host system publications.

Note: The SGI host system does not support the SDD. IRIX manages multipathing.

Checking the version of the IRIX operating system for the SGI host system

To confirm that you have the correct version of the IRIX operating system installed on your host system, type `# uname -Rs`.

Installing a fibre-channel adapter card for the SGI host system

This section lists the steps to install a fibre-channel adapter card for the SGI host system.

1. Install the adapter card in the host system.
2. Connect the cable to the storage unit port.
3. Restart the server.

Note: The IRIX driver is part of the IRIX operating system. There are no additional driver packages. SGI is responsible for support and distribution of the adapter when you use it in an SGI host system.

Verifying the installation of a fibre-channel adapter card for SGI

Use the following steps to verifying your installation of a fibre-channel adapter card for SGI.

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

Figure 80 on page 167 shows example results for the `# ls -l /hw/scsi_ctlr` command. In this example, the PCI host bus adapters are installed in slots 6 and 8 and the XIO host bus adapters are installed in slots 7 and 9.

```
# ls -l /hw/scsi_ctlr
lrw----- 1 root sys 50 Mar 27 10:31 6 -> /hw/module/1/slot/io7/xtalk_pci/pci/1/scsi_ctlr/0
lrw----- 1 root sys 48 Mar 27 10:31 7 -> /hw/module/1/slot/io2/pci_xio/pci/0/scsi_ctlr/0
lrw----- 1 root sys 50 Mar 27 10:31 8 -> /hw/module/1/slot/io6/xtalk_pci/pci/1/scsi_ctlr/0
lrw----- 1 root sys 48 Mar 27 10:31 9 -> /hw/module/1/slot/io2/pci_xio/pci/1/scsi_ctlr/0
```

Figure 80. Example of PCI bus slots

2. Type `scsiha -w /hw/scsi_ctlr/<N>/bus`, where `<N>` is the number of the PCI bus, to confirm the installation of the fibre-channel driver.

Configuring the fibre-channel adapter drivers for SGI

Use the commands listed in this section to configure fibre-channel adapter drivers for SGI.

1. Type `# cd /var/sysgen/master.d` to navigate to the necessary directory.
2. Type `# vi qlfc` to open the file.
3. Edit and configure the adapter driver for switched-vfabric or arbitrated loop topology. Set `qlfc_use_connection_mode` to 0 for arbitrated loop or to 1 for switched-fabric.

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

Use the commands in this section to install an optical cable for SGI in a switched-fabric topology.

1. Connect the fibre-channel host adapters from the SGI host system to an IBM 2109 or a Brocade switch.
2. Connect the storage unit fibre-channel host adapters to the switch.

Installing an optical cable for SGI in an arbitrated loop topology

Use the commands in this section to install an optical cable for SGI in an arbitrated loop topology.

1. Connect the fibre-channel host adapters from the SGI host system directly to the fibre-channel host adapters for the storage unit.
2. Connect the storage unit fibre-channel host adapters to the switch.

Confirming switch connectivity for SGI

Use the following commands to confirm switch connectivity for SGI.

1. Log in to the IBM or Brocade switch as administrator.
2. At the console prompt, type `switchshow`.
3. Confirm that each QLogic host adapter and storage unit host adapter has performed a fabric log in to the switch.

Figure 81 on page 168 shows example results for the `switchshow` command for an SGI host. In this example, ports 9 and 11 are used for testing other devices.

```

snj2109s16h1:os1> switchshow
switchName:      snj2109s16h1
switchType:      2.4
switchState:     Online
switchRole:      Principal
switchDomain:     1
switchId:        fffc01
switchWwn:       10:00:00:60:69:12:30:25
switchBeacon:    OFF
port 0: sw Online      F-Port 21:00:00:e0:8b:04:d5:39
port 1: sw Online      F-Port 21:00:00:e0:8b:04:ee:3a
port 2: sw Online      F-Port 21:00:00:e0:8b:03:8f:e6
port 3: sw Online      F-Port 21:00:00:e0:8b:04:13:3a
port 4: sw Online      F-Port 21:00:00:e0:8b:05:c4:07
port 5: sw Online      F-Port 21:00:00:e0:8b:05:0f:08
port 6: sw Online      F-Port 21:00:00:e0:8b:05:d2:07
port 7: sw Online      F-Port 21:00:00:e0:8b:05:08:08
port 8: sw Online      F-Port 21:00:00:e0:8b:04:e9:39
port 9: sw Online      L-Port 1 public
port 10: sw Online     F-Port 21:00:00:e0:8b:03:97:e6
port 11: sw Online     L-Port 1 public
port 12: sw Online     F-Port 21:00:00:e0:8b:05:d5:07
port 13: sw Online     F-Port 50:05:07:63:00:c3:91:74
port 14: sw Online     F-Port 21:00:00:e0:8b:05:07:08
port 15: sw Online     F-Port 50:05:07:63:00:c7:91:74

```

Figure 81. Example results for the switchshow command

Displaying zoning information for the switch

This section describes the command needed to display zoning information for the switch.

Use the **cfgShow** command to display the zoning information of the switch.

Figure 82 on page 169 shows a zoning example based on the unique worldwide port names of the host adapters and the storage unit fibre-channel ports. This is known as soft zoning. Each path to a storage unit LUN has a separate zone. For information on locating the worldwide port name for an SGI host, see “Locating the WWPN for an SGI host” on page 229.


```

osvl2109c:admin> cfgShow
Defined configuration:
cfg:  osvl_lab
      dc_1; w2k_1; MS6000Cluster; MSHPCluster; Tivoli; compaqzone1;
      compaqzone2; MS8500Cluster; AIX_ZONE; OSPL3Zone;
      MSCompacqCluster; SGI_Path1; SGI_Path2; NW;
      Netfyzone1; Netfyzone2; Netfyzone3; Netfyzone4

...
zone:  SGI_Path1
      SGI01; osp13_b1a2
zone:  SGI_Path2
      SGI02; osp13_b4a2
...
alias: SGI01    21:00:00:e0:8b:04:d5:39
alias: SGI02    21:00:00:e0:8b:04:ee:3a
alias: osp13_b1a2
           50:05:07:63:00:c3:91:74

alias: osp13_b4a2
           50:05:07:63:00:c7:91:74

...
Effective configuration:
cfg:  osvl_lab
...
zone:  SGI_Path1
      21:00:00:e0:8b:04:d5:39
      50:05:07:63:00:c3:91:74

zone:  SGI_Path2
      21:00:00:e0:8b:04:ee:3a
      50:05:07:63:00:c7:91:74

...

```

Figure 82. Example results for the `cfgShow` command

Confirming storage connectivity

The following sections tell you how to confirm storage connectivity for the following topologies:

- Switched fabric
- Arbitrated loop

Confirming storage connectivity for SGI in a switched-fabric topology

Confirming storage connectivity for SGI in a switched-fabric topology

This section contains an example of a switched fabric connection with the fibre-channel adapter cards. The following conditions apply to this example:

- The PCI-FC-1P-OPT adapter cards are on PCI busses 6 and 8.
- There are four LUNs defined on the storage unit.
- There are two connections from the switch to two host adapters on the storage unit.
- The switch is not zoned. Therefore, each LUN is available on each host adapter from each host bus adapter. This means that there are four instances of each LUN for each adapter card.

1. Type `scsiha -r {bus_number | device}` at the root of the shell prompt.
bus_number is the numerical value for the bus and device is numerical value of the device.
2. Type `scsiha -p {bus_number | device}`.
3. Use the `ioconfig -d -f /hw` command to assign logical controller numbers in the hardware graph to each of the newly discovered physical devices.
4. Use the `hinv -c disk` command to confirm that all four instances of each of the four LUNs exists.
5. Use the `ls -l /hw/rdisk` command to verify that you named the directory the same name as the worldwide node name of the storage unit.
6. Use the `ls -l /hw/rdisk` command to verify that directories for each LUN exist in the `/hw/rdisk/<wwnn>` directory, where <wwnn> is the worldwide node name of the device.
7. Type `cd /hw/rdisk/<wwnn>`, where <wwnn> is the worldwide node name of the device, to change to the `/hw/rdisk/<wwnn>` directory. You must execute the next command from within this directory.
8. Use the `ls -lR lun*` command to verify that the command created the character-special files in each LUN* directory. The asterisk (*) acts as a universal wildcard for any characters that might appear after the initial "LUN" characters.

Confirming storage connectivity for SGI in a fibre-channel arbitrated loop topology

Confirming storage connectivity for SGI in a fibre-channel arbitrated loop topology

This section contains an example of an arbitrated loop connection with the fibre-channel adapter cards. The following conditions apply to this example:

- The PCI-FC-1P-OPT adapter cards are on PCI busses 6 and 8.
 - There are four LUNs defined on the storage unit.
 - Each host adapter is connected to a host adapter in the storage unit.
 - Each host adapter maintains two instances of each LUN.
1. Type `scsiha -r {bus_number | device}` at the root of the shell prompt.
bus_number is the numerical value for the bus and device is the numerical value for the device.
 2. Type `scsiha -p {bus_number | device}`.
 3. Use the `ioconfig -d -f /hw` command to assign logical controller numbers in the hardware graph to each of the newly discovered physical devices.
 4. Use the `hinv -c disk` command to confirm that two instances of each of the four LUNs exist.
 5. Use the `ls -l /hw/disk` command to verify that the command created the appropriate block special files in `hw/disk` for each LUN.

Configuring the storage unit for host failover

This section describes how to configure the storage unit for host failover.

Configuring the storage unit for host failover

Configuring the storage unit for host failover

IRIX multipathing does not perform dynamic load-balancing; it only provides path failover. You can achieve static load-balancing if you distribute the primary paths evenly across the controllers when you set up the failover groups in the `/etc/failover.conf` file.

When a failover occurs to the alternate path, you must use the `/sbin/scsifo -d` command to manually return to the primary path.

When you have a host system with multiple host bus adapters or multiple paths to the storage, IRIX automatically chooses the first path it encounters in the hardware tree as the path to the storage. To prevent this from occurring when you perform a static load-balance, you must configure the `failover.conf` file to point the I/O operations to the correct path. This ensures that you assign a failover path or a backup path for all I/O operations.

Host failover using IRIX multipathing requires the use of the XLV volume manager or in the case of IRIX clusters, the XVM volume manager.

Confirming the availability of failover

Confirming the availability of failover.

1. Use the **chkconfig** command to determine if failover is enabled.
2. If failover is not enabled, turn it on and start the failover script or restart the server. Figure 83 shows an example of what commands to type if failover is disabled.

```
#chkconfig | grep failover
failover                off
#chkconfig failover on
#/etc/init.d/failover init
#Configuring Failover.
```

Figure 83. Example of commands to turn failover on

Making a connection through a switched-fabric topology

Making a connection through a switched-fabric topology

1. Edit the `/etc/failover.conf` file to add primary and secondary paths to the storage unit LUNs. “Confirming storage connectivity for SGI in a switched-fabric topology” on page 169 describes the paths of the character-special files. These files begin with the worldwide node name of the storage unit. Figure 84 on page 172 shows an example of an edited `/etc/failover.conf` file.

```

c6lun0      5005076300c09174/lun0/c6p30000 \
            5005076300c09174/lun0/c8p70000

c6lun1      5005076300c09174/lun1/c8p70000 \
            5005076300c09174/lun1/c6p30000

c6lun2      5005076300c09174/lun2/c6p30000 \
            5005076300c09174/lun2/c8p70000

c6lun3      5005076300c09174/lun3/c8p70000 \
            5005076300c09174/lun3/c6p30000

```

Figure 84. Example of an edited `/etc/failover.conf` file

2. Use the IRIX `/sbin/scsifo -d` command to confirm the failover configuration.

Making a connection through an arbitrated-loop topology

Making a connection through an arbitrated-loop topology

1. Edit the `/etc/failover.conf` file to add primary and secondary paths to storage unit LUNs. Figure 85 shows an example of an edited `/etc/failover.conf` file that adds paths to the storage unit LUNs.

```

GroupA sc3d0111 sc4d0111
GroupB sc3d0117 sc4d0117

```

Figure 85. Example of an edited `/etc/failover.conf` file for an arbitrated loop connection

2. Use the IRIX `/sbin/scsifo -d` to confirm the proper failover configuration.

Switching I/O operations between the primary and secondary paths

Switching I/O operations between the primary and secondary paths

Use the `sbin/scsifo -s` command to manually switch from the primary path to the alternate path. With this command, you can also specify the path from where you want to switch within a failover group. For example, you might want to perform this action to switch back to the initial primary path after a path failure.

Configuring storage

The following sections describe how to configure storage in either a switched-fabric or arbitrated-loop topology.

Configuration considerations

This section provides configuration considerations.

Note: The `/hw` directory is used to build the hardware graph. The hardware graph represents the collection of all significant hardware that is connected to a system. Do not specify the `/hw` entries in commands that pertain to disk devices. Use the traditional `/dev/rdisk` entries instead. The following example shows a naming convention for the disk devices.

```
#/usr/bin/fx -x -d /dev/rdisk/5005076300c003b4/1un1vo1/c6p30000
```

Configuring storage in a switched fabric topology

Configuring storage in a switched fabric topology

Figure 86 on page 174 and Figure 87 on page 175 shows the standard IRIX storage configuration utilities that you can use to partition and format the storage unit LUNs and to create and mount the file systems. The commands are in bold type.

Note: You must be a root user to use these commands.

```

#/usr/bin/fx -x -d /dev/rdisk/5005076300c003b4/1un1vol/c6p30000
fx version 6.5, Jul  7, 2000
...opening /dev/rdisk/5005076300c003b4/1un1vol/c3p1000

...drive selftest...OK
Scsi drive type == IBM      2105F20      1206

----- please choose one (? for help, .. to quit this menu)-----
[exi]t          [d]ebug/          [l]abel/          [a]uto
[b]adblock/     [ex]ercise/       [r]epartition/

fx> repartition

----- partitions-----
part type      blocks          Megabytes  (base+size)
 7: xlv         4096 + 71798784      2 + 35058
 8: volhdr      0 + 4096             0 + 2
10: volume      0 + 71802880        0 + 35060

capacity is 71802880 blocks

----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive     [o]ptiondrive       [e]xpert
[us]rrootdrive  [re]size
fx/repartition> o

fx/repartition/optiondrive: type of data partition = (xfs) xfs
Warning: you will need to re-install all software and restore user data
from backups after changing the partition layout. Changing partitions
will cause all data on the drive to be lost. Be sure you have the drive
backed up if it contains any user data. Continue? yes

----- partitions-----
part type      blocks          Megabytes  (base+size)
 7: xfs         4096 + 71798784      2 + 35058
 8: volhdr      0 + 4096             0 + 2
10: volume      0 + 71802880        0 + 35060

capacity is 71802880 blocks

----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive     [o]ptiondrive       [e]xpert
[us]rrootdrive  [re]size
fx/repartition> ..

----- please choose one (? for help, .. to quit this menu)-----
[exi]t          [d]ebug/          [l]abel/          [a]uto
[b]adblock/     [ex]ercise/       [r]epartition/

fx>exit

```

Figure 86. Example commands for the IRIX switched fabric storage configuration utility

```

# xlv_make
xlv_make> vol Sharks0
Sharks0
xlv_make> data
Sharks0.data
xlv_make> plex
Sharks0.data.0
xlv_make> ve -force "/dev/dsk/5005076300c003b4/1un1s7/c6p3s0000"
xlv_make> quit

# ls -l /dev/xlv
total 0
brw----- 1 root sys 192, 10 Mar 12 16:06 Sharks0

# mkfs -t xfs /dev/xlv/Sharks0
meta-data=/dev/xlv/Sharks0 isize=256 agcount=8, agsize=27403 blks
data = bsize=4096 blocks=219218, imaxpct=25
= sunit=0 swidth=0 blks, unwritten=1
naming =version 1 bsize=4096
log =internal log bsize=4096 blocks=1168
realtime =none extsz=65536 blocks=0, rtextents=0

# mkdir /lv1_mount
# mount -t xfs /dev/xlv/Sharks0 /lv1_mount
# df -k

```

Filesystem	Type	kbytes	use	avail	%use	Mounted on
/dev/root	xfs	1961580	1750112	211468	90	/
/dev/xlv/Sharks0	xfs	872200	144	872056	1	/lv1_mount

```

#

```

Figure 87. Example commands for the IRIX switched fabric storage configuration utility, part 2

Configuring storage in an arbitrated loop topology

This section provides information for configuring storage in an arbitrated loop topology.

Figure 88 on page 176 and Figure 89 on page 177 shows the standard IRIX storage configuration utilities that you can use to partition and format the storage unit LUNs and to create and mount file systems. The commands are in bold type.

```

# fx -x -d /dev/rdsk/dks3d0117s0

...drive selftest...OK
Scsi drive type == IBM      2105F20      1206

----- please choose one (? for help, .. to quit this menu)-----
[exi]t          [d]ebug/          [l]abel/          [a]uto
[b]adblock/     [exe]rcise/       [r]epartition/

fx> repartition

----- partitions-----
part type      blocks          Megabytes  (base+size)
 7: xlv        4096 + 71798784      2 + 35058
 8: volhdr     0 + 4096         0 + 2
10: volume     0 + 71802880      0 + 35060

capacity is 71802880 blocks

----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive     [o]ptiondrive    [e]xpert
[u]srrootdrive  [re]size
fx/repartition> o

fx/repartition/optiondrive: type of data partition = (xfs) xfs
Warning: you will need to re-install all software and restore user data
from backups after changing the partition layout. Changing partitions
will cause all data on the drive to be lost. Be sure you have the drive
backed up if it contains any user data. Continue? yes

----- partitions-----
part type      blocks          Megabytes  (base+size)
 7: xfs        4096 + 71798784      2 + 35058
 8: volhdr     0 + 4096         0 + 2
10: volume     0 + 71802880      0 + 35060

capacity is 71802880 blocks
----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive     [o]ptiondrive    [e]xpert
[u]srrootdrive  [re]size
fx/repartition> ..

----- please choose one (? for help, .. to quit this menu)-----
[exi]t          [d]ebug/          [l]abel/          [a]uto
[b]adblock/     [exe]rcise/       [r]epartition/

fx>exit

```

Figure 88. Example commands for the IRIX arbitrated loop storage configuration utility


```

# xlv_make
xlv_make> vol Sharks0
Sharks0
xlv_make> data
Sharks0.data
xlv_make> plex
Sharks0.data.0
xlv_make> ve -force "/dev/rdisk/dks3d0117s0"
xlv_make> quit

# ls -l /dev/xlv
total 0
brw----- 1 root    sys      192, 10 Mar 12 16:06 Sharks0

# mkfs -t xfs /dev/xlv/Sharks0
meta-data=/dev/xlv/Sharks0    isize=256    agcount=8, agsize=27403 blks
data      =                   bsize=4096    blocks=219218, imaxpct=25
          =                   sunit=0       swidth=0 blks, unwritten=1
naming    =version 1          bsize=4096
log       =internal log      bsize=4096    blocks=1168
realtime  =none              extsz=65536    blocks=0, rtextents=0

# mkdir /lv1_mount
# mount -t xfs /dev/xlv/Sharks0 /lv1_mount
# df -k

```

Filesystem	Type	kbytes	use	avail	%use	Mounted on
/dev/root	xfs	1961580	1750112	211468	90	/
/dev/xlv/Sharks0	xfs	872200	144	872056	1	/lv1_mount

```

#

```

Figure 89. Example commands for the IRIX arbitrated loop storage configuration utility, part 2

Chapter 18. Attaching to a Sun host

This chapter provides instructions, requirements, and considerations for attaching a Sun host system to a storage unit.

Attaching with fibre-channel adapters

This section provides information for attaching a storage unit to a Sun host system with fibre-channel adapters. It also describes how to install, download, and configure the adapter cards.

This section also tells you how to change the Sun system kernel. Before you start, you must meet the attachment requirements listed in “Fibre-channel attachment requirements for Sun.”

Supported fibre-channel adapters for Sun

This section lists the fibre-channel adapter cards that are supported for the Sun host system.

The following fibre-channel adapter cards are supported for the Sun host system:

- Emulex LP9002L adapter card
- Emulex LP9002DC adapter card
- Emulex LP9002S adapter card
- Emulex LP9402DC adapter card
- Emulex LP9802 adapter card
- Emulex LP10000 adapter card
- Emulex LP 10000DC adapter card
- AMCC FCE-1473 adapter card
- AMCC FCE-6460 adapter card
- AMCC FCX-6562 adapter card
- AMCC FCX2-6562 adapter card
- QLogic QLA2310FL adapter card
- QLogic QLA2340FL adapter card
- QLogic QLA2342FL adapter card
- Sun Adapter 6727A adapter card
- Sun Adapter 6757A adapter card
- Sun Adapter 6767A adapter card (SG-XPCI1FC-QF2)
- Sun Adapter 6768A adapter card (SG-XPCI2FC-QF2)
- Sun Adapter 6799A adapter card

Note: For a list of open-systems hosts, operating systems, adapters and switches that IBM does not support, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Fibre-channel attachment requirements for Sun

This section lists the requirements for attaching the storage unit to your host system.

This section lists the requirements for attaching the storage unit to your host system:

- Ensure that there are enough fibre-channel adapters installed in the server to manage the total LUNs you want to attach.
- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *User's Guide* is on the TotalStorage CD that you receive with the storage unit.
- Solaris 8 and Solaris 9 require patches to ensure that the host and the storage unit function correctly. Table 11 provides the minimum revision level that is required for each Solaris patch ID.

Table 11. Solaris 8 minimum revision level patches for fibre-channel

Solaris 8
108528-03 kernel update
109524-02 ssd driver
109657-01 isp driver
108974-03 sd, uata drivers
Not applicable

- Review device driver installation documents and configuration utility documents for additional Solaris patches that you might need.
- For details about the release level for your operating system, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Either you or an IBM service support representative must perform the following tasks to install and configure a storage unit:

1. Install the storage unit.
2. Define the fibre-channel host system with the worldwide port name identifiers. For the list of worldwide port names, see "Locating the worldwide port name (WWPN)," on page 225.
3. Define the fibre-port configuration if you did not do it when you installed the storage unit or fibre-channel adapters.
4. Configure the host system for the storage unit by using the instructions in your host system publications.

Note: SDD does not support the Sun host system in a clustering environment. To have failover protection on an open system, SDD requires a minimum of two fibre-channel adapters. The maximum number of fibre-channel adapters supported is 16 for a total of 16 fibre-channel ports.

You can use the SDD for the following Sun host system modes:

- Sun host hardware platforms with 32- or 64-bit mode capabilities that run Solaris 8 or Solaris 9.

Installing the Emulex adapter card for a Sun host system

This section tells you how to attach a storage unit to a Sun host system with an Emulex adapter card.

For the recommended host adapter parameter settings, see *Host Bus Adapter FC Support* at <http://www-1.ibm.com/servers/storage/disk/ds8000>.

Note: The Emulex LP9402DC adapter card is a 3.3 volt card only and requires a 3.3 volt PCI slot.

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

- Target
 - Public initiator
 - Private initiator
 - Target and public initiator
 - Target and private initiator
1. Record the IEEE number that is printed on the card. You can use the IEEE number to determine the WWPNN.
 2. Refer to the installation instructions provided by your host adapter vendor for their specific installation instructions. See <http://www.emulex.com> for the latest documentation.

Downloading the Emulex adapter driver for a Sun host system

This section provides instructions to download the current Emulex fibre-channel adapter driver.

1. Restart your host system.
2. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
3. Click **Interoperability matrix**.
4. Click **DS8000 interoperability matrix**.
5. Find the section for the current version of the driver and firmware that you want.
6. Go to <http://www.emulex.com>.
7. Click **Driver, downloads, and documentation** from the left navigation pane. Click **OEM software and documentation**.
8. Select **IBM** from the drop down vendor list.
9. Click the link for the adapter that corresponds to the firmware, driver, and documentation that you need to install and download the adapter driver.

Installing the Emulex adapter driver for a Sun host system

This section provides instructions to install the Emulex fibre-channel adapter driver.

1. Go to www.emulex.com.
2. From the left navigation pane, click **Drivers, Software, and Manuals**.
3. Click **IBM**.
4. Determine the driver you want, and click the link for the adapter driver.
5. Determine the operating system for the driver that you want to download and click **Version x-x.xxxx**, where x-x.xxxx is the version number for the adapter driver.

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

Installing the AMCC PCI adapter card for Sun

This section tells you how to attach a storage unit to a Sun host system running Solaris 8 or 9 with the AMCC PCI adapter card.

1. Shutdown your system.

2. Insert the AMCC PCI adapter into a PCI slot and connect the adapter to your fibre-channel disk.
3. Restart the system.

Downloading the current AMCC PCI adapter driver for Sun

This section tells you how to download the AMCC PCI fibre-channel adapter driver.

1. Restart your host system.
2. Go to <http://www.amcc.com>
3. From the navigation menu at the top of the page, click **Drivers/Downloads**.
4. Select the type of driver that you are interested in and follow the instructions from the web site.
5. If you downloaded the driver file from a Sun host system, go to “Installing the AMCC PCI adapter card for Sun” on page 181. If you downloaded the driver file from a non-Sun host system, transfer the drive file to a Sun host system, and then go to “Installing the AMCC PCI adapter card for Sun” on page 181.

Installing the AMCC PCI adapter driver for Sun

This topic provides instructions to install the AMCC PCI adapter drivers.

1. Go to <http://www.amcc.com>
2. From the navigation menu at the top of the page, click **Drivers/Downloads**.
3. Select the type of driver that you are interested in and follow the instructions from the web site.
4. Follow the instructions in the readme.txt file to install the AMCC PCI adapter card.
5. Update the parameter list and restart the host system. See Table 13 on page 187 for the parameters and recommended settings.

Installing the AMCC SBUS adapter card for Sun

This section tells you how to attach a storage unit to a Sun host system running Solaris 8 or 9 with the AMCC SBUS adapter card.

1. Shutdown your system.
2. Insert the AMCC SBUS adapter into a empty SBUS slot and connect the adapter to your fibre-channel disk.
3. Restart the system.

Downloading the current AMCC SBUS adapter driver for Sun

This section provides instructions to download the AMCC SBUS adapter driver.

1. Restart your host system.
2. Go to <http://www.amcc.com>
3. From the navigation menu at the top of the page, click **Drivers/Downloads**.
4. Select the type of driver that you are interested in and follow the instructions from the web site.
5. Install the AMCC PCI adapter driver. If you download the driver from a non-Sun host system, transfer the driver file to a Sun host system before installing it.

Installing the AMCC SBUS adapter driver for Sun

This section provides instructions to install the AMCC SBUS adapter driver.

1. Go to <http://www.amcc.com>

2. From the navigation menu at the top of the page, click **Drivers/Downloads**.
3. Select the type of driver that you are interested in and follow the instructions from the web site.
4. Follow the instructions in the readme.txt file to install the AMCC SBUS adapter card.
5. Update the parameter list and restart the host system. See Table 13 on page 187 for the parameters and recommended settings.

Installing the QLogic adapter card in the Sun host

This section tells you how to attach a storage unit to a Sun host system running Solaris 8 or 9 with the QLogic adapter cards.

1. Shutdown your system.
2. Insert the QLogic adapter into a PCI or SBUS slot for the QLogic adapter in the system, and connect the adapter to your fibre-channel disk.
3. Restart the system.

Downloading the QLogic adapter driver for Sun

This section provides instructions to download the current adapter driver onto the QLogic adapter card.

1. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
2. Click **Interoperability matrix**.
3. Click **DS8000 interoperability matrix**.
4. Find the section for the current version of the driver and firmware and driver you want.
5. Go to <http://www.qlogic.com>.
6. Click **Drivers/Software** from the **Quick Links** list.
7. Choose **IBM** from the **OEM and ISV approved drivers and firmware** section.
8. Navigate to the part of the page for the adapter you want to download, then click **download**.

Installing the QLogic adapter driver package for Sun

This section provides instructions to install the QLogic host adapter driver package.

1. Obtain a copy of the package image file, q1a23xx.Z by performing the steps in "Downloading the QLogic adapter driver for Sun."
2. Log in to your Solaris system as root.
3. Type `cd/tmp` to change the directory to a temporary directory.
4. Type the following command to copy the compressed package file to the /tmp directory:
`cp <source directory>/q1a23xx.Z`
 where <source directory> is the directory where the compressed package file resides after downloading the driver from the QLogic Web site.
5. Type the following command to uncompress the package file:
`uncompress q1a23xx.Z`

Note: You must use an upper case Z when you specify the file name.

6. Transfer the QLogic driver package to a diskette before you use the package to install the adapter drivers. Use the following procedure to create a diskette that contains the QLogic driver package:
 - a. Type `/etc/init.d/volmgt stop` to stop the volume manager.
 - b. Type `dd if=q1a23xx of=/dev/rdiskette0` to transfer the image file.

Configuring host device drivers

The following instructions explain how to update device driver configuration files on the Sun host to enable access to target and LUN pairs that are configured on the storage unit.

Configuring host device drivers for Sun

The following instructions explain how to update device driver configuration files on the Sun host to enable access to target and LUN pairs that are configured on the storage unit.

Perform the following steps to update the Solaris driver configuration file. This procedure gives you access to target and LUN pairs that are configured on the storage unit:

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

1. Change to the directory by typing: `cd /kernel/drv`
2. Backup the sd.conf file in this subdirectory.
3. Edit the sd.conf file to add support for the target and LUN pairs that are configured on the host system.

Note: Do not add duplicate target and LUN pairs.

Figure 90 shows the lines that you must add to the file to access LUNs 0 - 49 on target 0 for fibre-channel.

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

Figure 90. Example of sd.conf file entries for fibre-channel

Figure 91 on page 185 shows the start lpfc auto-generated configuration.

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

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

Figure 91. Example of a start lpfc auto-generated configuration

4. Type either:

- a. `reboot -- -r` from the Open Windows window to shutdown and restart the Sun host system with the kernel reconfiguration option
- b. `boot -r` from the OK prompt after you shutdown

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

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

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

Parameter settings for the Emulex adapters for the Sun host system

You can use these recommended configuration settings for your Emulex adapters on a Sun host system.

Table 12 on page 186 provides configuration settings that are recommended for the Emulex adapters.

Table 12. Recommended configuration file parameters for the Emulex LP9002DC, LP9002L, LP9002S, LP9402DC and LP9802 adapters

Parameters	Recommended settings
automap	1: Default. SCSI IDs for all FCP nodes without persistent bindings are automatically generated. If new FCP devices are added to the network when the system is down, there is no guarantee that these SCSI IDs will remain the same when the system is restarted. If one of the FCP binding methods is specified, then automap devices use the same mapping method to preserve SCSI IDs between link down and link up. If no bindings are specified, a value of 1 forces WWNN binding, a value of 2 forces WWPN binding, and a value of 3 forces DID binding. If automap is 0, only devices with persistent bindings are recognized by the system.
fcp-on	1: Default. Turn on FCP.
lun-queue-depth	30: The default value that the driver uses to limit the number of outstanding commands per FCP LUN. This value is global, affecting each LUN recognized by the driver, but can be overridden on a per-LUN basis. You might have to configure RAID using the per-LUN tunable throttles.
no-device-delay	0: Default. Implies no delay whatsoever. 1: Recommended. 2: Setting a long delay value might permit I/O operations to build up, each with a pending time-out, which could result in the exhaustion of critical Solaris kernel resources. In this case, you might see a fatal message such as, PANIC: Timeout table overflow.
network-on	0: Default. Recommended for fabric. Do not turn on IP networking. 1: Turn on IP networking.
scan-down	0: Recommended. Causes the driver to use an inverted ALPA map, effectively scanning ALPAs from high to low as specified in the FC-AL annex. 2: Arbitrated loop topology.
tgt-queue-depth	0: Recommended. The default value that the driver uses to limit the number of outstanding commands per FCP target. This value is global, affecting each target recognized by the driver, but can be overridden on a per-target basis. You might have to configure RAID using the per-target tunable throttles.

Table 12. Recommended configuration file parameters for the Emulex LP9002DC, LP9002L, LP9002S, LP9402DC and LP9802 adapters (continued)

Parameters	Recommended settings
topology	2: Recommended for fabric. Point-to-point topology only. 4: Recommended for nonfabric. Arbitrated-loop topology only.
xmt-queue-size	256: Default. Size of the transmit queue for mbufs (128 - 10240).
zone-rscn	0: Default. 1: Recommended for fabric. Check the name server for RSCNs. Setting zone-rscn to 1 causes the driver to check with the name server to see if an N_Port ID received from an RSCN applies. If soft zoning is used with Brocade fabrics, this parameter should be set to 1.

Parameter settings for the AMCC adapters for the Sun host system

You can use these recommended settings for your AMCC adapters on a Sun host system.

Table 13 provides configuration settings that are recommended for the supported adapters. For the most current information about fibre-channel adapter parameter settings, see:

<http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do>

Table 13. Recommended configuration file parameters for a AMCC FCX-6562, AMCC FCX2-6562, AMCC FCE-6460, or a AMCC FCE-1473 adapter

Parameters	Recommended settings
FcEngHeartbeatInterval	5: Default. When the AMCC adapter or driver detects that the fibre-channel link is up (and there is no I/O activity), it sends a test frame (or heartbeat) to itself to verify link integrity. The test frame is sent at the interval specified by this parameter. If the test frame does not complete, there might be a link problem. In this situation, the driver initiates error recovery to reestablish a good link. A value of 0 disables the heartbeat.
FcLinkUpRecoveryTime	1000: Default. Delay (msec) after the link is up before port discovery begins, allowing the link to stabilize and protecting against a possible I/O surge. This timer is reset every time the link comes up. The default value is adequate for most configurations.

Table 13. Recommended configuration file parameters for a AMCC FCX-6562, AMCC FCX2-6562, AMCC FCE-6460, or a AMCC FCE-1473 adapter (continued)

Parameters	Recommended settings
BusyRetryDelay	5000: Default. Delay (msec) before retrying after receipt of an I/O operation with a SCSI Busy status from a target. The number of retries is based on the Solaris retry count associated with the I/O operation.
FailoverDelay	30: Delay (seconds) before failing all I/O operations for an offline target. If the delay timer expires, all I/O operations for the failed target is returned to the application. A zero value disables failover.
TimeoutResetEnable	0: False. Boolean parameter for enabling SCSI target resets for timed-out I/O operations. When the timer expires (usually 60 seconds, as specified by the upper layers), the driver issues a target reset to attempt to clear the device (which might be either too busy to respond or stuck).
QfullRetryCount	5: Default. Number of times an I/O operation is retried due to receipt of a SCSI queue full status from a target. The delay between retries is based on the QfullRetryDelay parameter.
QfullRetryDelay	5000: Default. Delay (msec) before retrying after receipt of an I/O operation with a SCSI queue full status from a target. The number of retries is based on the QfullRetryCount parameter.
LunRecoveryInterval	50: Default. Sets the LUN I/O recovery interval (in msec) after the driver reconnects to a disk. It is a global parameter affecting all targets, and determines how long the driver waits after a port is discovered until sending I/O operations to that port. Some devices might require more time to flush an I/O operation that was in progress prior to a link going down; if this is the case, increase the value of this parameter.

Table 13. Recommended configuration file parameters for a AMCC FCX-6562, AMCC FCX2-6562, AMCC FCE-6460, or a AMCC FCE-1473 adapter (continued)

Parameters	Recommended settings
FcLinkSpeed	<p>3: Default.</p> <p>Specifies the desired fibre-channel link speed as follows:</p> <ul style="list-style-type: none"> • 0: default to SEEPROM setting • 1: force 1 gigabit per second • 2: force 2 gigabit per second • 3: auto negotiate link speed
JniCreationDelay	<p>5: Default.</p> <p>Delay (seconds) after driver creation to allow the network to stabilize, discover ports, and build the driver database. Increase this value if targets are being discovered too late in the restart process.</p>
FlogiRetryCount	<p>3: Default.</p> <p>Total number of fabric login (FLOGI) attempts before giving up logging in to a switch. Failure prevents participation on a fabric topology.</p>
FcFlogiTimeout	<p>10: Default.</p> <p>Specifies the amount of time (in seconds) that the driver waits for a fabric login (FLOGI) to be accepted. The value should be increased only if the switch to which the host adapter is connected requires more time to respond to a FLOGI. The number of retries is configured with the FlogiRetryCount parameter.</p>
PlogiRetryCount	<p>5: Default.</p> <p>Total number of port login (PLOGI) attempts before failing to log in to a SCSI target.</p>
PlogiControlSeconds	<p>30: Default.</p> <p>Defines the number of seconds that the driver waits for a successful port login (PLOGI) attempt. The maximum number of attempts is defined by the PlogiRetryCount parameter. Some devices might take longer to respond to PLOGIs; if this is the case, increase the value of this parameter.</p>
FcEmldEngTcbCount	<p>1789: Default.</p> <p>Total number of concurrent exchanges (also called transfer control blocks) that the adapter allows. To optimize performance, set this parameter to match the memory capacity of the hardware.</p>

Table 13. Recommended configuration file parameters for a AMCC FCX-6562, AMCC FCX2-6562, AMCC FCE-6460, or a AMCC FCE-1473 adapter (continued)

Parameters	Recommended settings
TargetOfflineEnable	Recommended setting: 0. Function: Flag for enabling target offline when target port fails. 1: Default. (Enable).

Parameter settings for the QLogic QLA23xxF adapters

You can use these recommended configuration settings for your QLogic adapters on a Sun host system.

Table 14 provides configuration settings that are recommended for the QLogic QLA2310F, QLA2340, and QLA2342 adapters. These settings are for a storage unit that is attached directly to or through a fabric switch. For the most current information about fibre-channel adapter parameter settings, see:

<http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do>.

Table 14. Recommended configuration file parameters for the QLogic QLA2310F, QLA2340, and QLA2342 adapters with driver level 4.03

Parameters	Recommended settings
hba0-max-frame-length	=2048;
hba0-max-iocb-allocation	=256;
hba0-execution-throttle	=31;
hba0-login-timeout	=4;
hba0-login-retry-count	=1;
hba0-fabric-retry-count	=10;
hba0-adapter-hard-loop-ID	=0;
hba0-enable or disable-64bit PCI DMA addressing	<ul style="list-style-type: none"> =0: disabled =1: enabled
hba0-enable or /disable-LIP-reset	=0: disabled
hba0-enable or disable-LIP-full-login	=1: enabled
hba0-enable or disable-target-reset	=0: disabled
hba0-reset-delay	=5;
hba0-port-down-retry-count	=30;
hba0-link-down-error	=1: enabled
hba0-loop-down-timeout	=60;
hba0-connection-options	<ul style="list-style-type: none"> =0: loop only (for direct attachment) =1: point-to-point only (for fabric attachment)
hba0-device-configuration-mode	=1: use port name
hba0-fc-tape	<ul style="list-style-type: none"> =0: disabled =1: enabled

Table 14. Recommended configuration file parameters for the QLogic QLA2310F, QLA2340, and QLA2342 adapters with driver level 4.03 (continued)

Parameters	Recommended settings
hba0-fc-data-rate	=2: Auto-negotiate
hba0-command-completion-option	=0: response queue
persistent binding only option	=0: reports the discovery of persistent binded and non-bindded devices to the operating system

Parameter settings for the QLogic QLA23xx adapters for San Surf configuration (4.06+ driver)

The following table provides the parameter settings for the QLogic QLA23xx adapters for storage units with the San Surf configuration (4.06+ driver).

Table 15 provides the QLogic QLA23xx adapter settings for storage units with the 4.06+ driver of the San Surf configuration.

Table 15. Parameter settings for the QLogic QLA23xx host adapters for San Surf Configuration (4.06+)

Adapter parameter	Recommended settings
Host Parameters	
Execution Throttle	16 (default)
Frame Length	2048
Port Down Retry Count	30
Connection Options	<ul style="list-style-type: none"> • Loop Only • Loop Retrieved • Point-to-Point
Enable Fibre Channel Tape Support	Enable
Enable Fibre Channel Confirm	Enable
Data Rate	2 Auto Negotiate
Advanced Host Parameters (See Note 1)	
Persistent Plus New	Enable
Persistent Target Only	Disable
Persistent Bind HBA	Disable
Fast Error Reporting	Disable
Link Down Error	Enable
Extended Logging	Disable
Enable Target Reset	Enable
Enable Fast Command Posting	Disable
Enable 4Gb Addressing	Disable
Maximum LUNs per Target	8 (default)
Loop Down Timeout	60 (default)
Advanced Firmware Parameters (See Note 1)	
Enable Adapter Hard Loop ID	Disable

Table 15. Parameter settings for the QLogic QLA23xx host adapters for San Surf Configuration (4.06+) (continued)

Adapter parameter	Recommended settings
Enable LIP Full Login	Enable
Enable Class 2 Service	Disable
Enable ACKO	Disable
Enable Read Transfer Ready	Disable
Hard Loop ID	<default>
Login Retry Count	8
Interrupt Delay Timer	0
Wizard (for binding) (See Note 2)	
Note: <ol style="list-style-type: none"> 1. All of the parameters in the Advanced Host Parameters and Advanced Firmware Parameters categories use the default parameter settings. 2. San Surf uses Wizard to update the binding configuration information in the qlaxxxx.conf file. See Figure 92 for an example of the binding inserts. 	

Figure 92 shows an example qlaxxxx.conf file that contains binding inserts at the end of the file.

```
hba3-SCSI-target-id-0-fibre-channel-node-name="5005076300c0863a";
hba3-SCSI-target-id-0-fibre-channel-port-name="5005076300cf863a";

hba3-SCSI-target-id-1-fibre-channel-node-name="5005076300c0863a";
hba3-SCSI-target-id-1-fibre-channel-port-name="5005076300cb863a";

hba2-SCSI-target-id-0-fibre-channel-node-name="5005076300c0863a";
hba2-SCSI-target-id-0-fibre-channel-port-name="5005076300cf863a";

hba2-SCSI-target-id-1-fibre-channel-node-name="5005076300c0863a";
hba2-SCSI-target-id-1-fibre-channel-port-name="5005076300cb863a";

hba1-SCSI-target-id-0-fibre-channel-node-name="5005076300c0863a";
hba1-SCSI-target-id-0-fibre-channel-port-name="5005076300c1863a";

hba1-SCSI-target-id-1-fibre-channel-node-name="5005076300c0863a";
hba1-SCSI-target-id-1-fibre-channel-port-name="5005076300c5863a";

hba0-SCSI-target-id-0-fibre-channel-node-name="5005076300c0863a";
hba0-SCSI-target-id-0-fibre-channel-port-name="5005076300c5863a";

hba0-SCSI-target-id-1-fibre-channel-node-name="5005076300c0863a";
hba0-SCSI-target-id-1-fibre-channel-port-name="5005076300c1863a";
```

Figure 92. Example binding inserts for qlaxxxx.conf

Setting the Sun host system parameters

The following sections contain the procedures to set the Sun host system parameters for optimum performance on the storage unit with the following adapters:

- AMCC
- Emulex
- QLogic

Setting parameters for AMCC adapters

This section provides instructions for setting parameters for AMCC adapters.

1. Type `cd /etc` to change to the `/etc` subdirectory.
2. Backup the system file in the subdirectory.
3. Edit the system file and set the following parameters for servers with configurations that use AMCC adapters:

sd_max_throttle

This `sd_max_throttle` parameter specifies the maximum number of commands that the `sd` driver can queue to the host adapter driver. The default value is 256, but you must set the parameter to a value less than or equal to a maximum queue depth for each LUN connected. Determine the value by using the following formula: $256 \div (\text{LUNs per adapter})$ where LUNs per adapter is the largest number of LUNs assigned to a single adapter.

To set the `sd_max_throttle` parameter for the storage unit LUNs in this example, you would add the following line to the `/etc/system` file: `set sd:sd_max_throttle=5`

sd_io_time

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

sd_retry_count

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

maxphys

This parameter specifies the maximum number of bytes that you can transfer for each SCSI transaction. The default value is 126976 (124 KB). If the I/O block size that you requested exceeds the default value, the request is broken into more than one request. You should tune the value for the application requirements. For maximum bandwidth, set the `maxphys` parameter by adding the following line to the `/etc/system` file: `set maxphys=1048576 (1 MB)`

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

Setting parameters for Emulex or QLogic adapters

This section provides instructions for setting parameters for Emulex or QLogic adapters.

1. Type `cd /etc` to change to the `/etc` subdirectory.
2. Backup the system file in the subdirectory.
3. Edit the system file and set the following parameters for servers with configurations that only use Emulex or QLogic adapters.

sd_max_throttle

This `sd_max_throttle` parameter specifies the maximum number of commands that the `sd` driver can queue to the host adapter driver. The default value is 256, but you must set the parameter to a value less than or equal to a maximum queue depth for each connected LUN.

Determine the value by using the following formula: $256 \div (\text{LUNs per adapter})$ where LUNs per adapter is the largest number of LUNs that are assigned to a single adapter.

To set the `sd_max_throttle` parameter for the storage unit LUNs in this example, add the following line to the `/etc/system` file: set `sd:sd_max_throttle=5`

sd_io_time

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

sd_retry_count

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

maxphys

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

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

Installing the Subsystem Device Driver

This section provides instructions for installing the Subsystem Device Driver.

For procedures about how to install the Subsystem Device Driver (SDD), see *IBM TotalStorage Multipath Subsystem Device Driver User's Guide*. This publication describes how to use the SDD on open-systems hosts to enhance performance and availability on the storage unit. The SDD creates redundant paths for shared logical unit numbers and permits applications to run without interruption when path errors occur. It balances the workload across paths, and it transparently integrates with applications. For information about the SDD, go to <http://www-1.ibm.com/servers/storage/support/software/sdd.html>.

Attaching a storage unit to a Sun host using Storage Traffic Manager System

This section provides instructions for attaching a storage unit to a Sun host using Storage Traffic Manager System.

The Storage Traffic Manager System (STMS) is a multipathing mechanism for native Sun systems. Using the DS CLI, you can specify **Sun - Solaris** as port profile. Using the Storage Manager GUI, you can specify **Sun servers (Solaris) (Sun)** as host system when attaching the storage unit to a host operating with STMS.

See the Sun Web site at <http://www.sun.com> for up-to-date documentation.

Configuring the Sun STMS host settings

This section provides instructions for configuring the Sun STMS host settings.

1. Type `cd /etc` to change to the `/etc` subdirectory.
2. Backup the system file in the subdirectory.
3. Edit the system file and set the following parameters for servers with configuration that only use Emulex or QLogic adapters:

ssd_io_time

This parameter specifies the time-out value for disk operations. Add the following line to the `/etc/system` file to set the `ssd_io_time` parameter for the storage unit LUNs:

```
set ssd:ssd_io_time=0x78
```

ssd_max_throttle

This `ssd_max_throttle` parameter specifies the maximum number of commands that the SSD driver can queue to the host adapter driver. The default value is 256, but you must set the parameter to a value less than or equal to a maximum queue depth for each connected LUN. Determine the value by using the following formula:

$$256 \div (\text{LUNs per adapter})$$

where LUNs per adapter is the largest number of LUNs assigned to a single adapter.

To set the `ssd_max_throttle` parameter for the storage unit LUNs in this example, add the following line to the `/etc/system` file:

```
set ssd:ssd_max_throttle=5
```

Attaching a storage unit to a Sun host using Sun Cluster

This section provides instructions for attaching a storage unit to a Sun host using Sun Cluster.

See the Sun Web site at <http://www.sun.com> for up-to-date documentation.

Attention: Due to a problem found in two-node SunCluster implementations on a storage unit, the following restriction must be observed. When setting up a two-node SunCluster 3.x environment with MPxIO for multipathing, it is mandatory that 1:1 zoning be used; that is, each HBA must be in a zone that contains only 1 HBA and only 1 storage unit target adapter. To ensure that this relationship is maintained, create host nickname definitions using the GUI to enforce "port masking" that allows the host WWPN to login *only* to the desired storage unit target adapter.

This zoning and port masking must be implemented prior to assigning any storage unit LUNs to the SunCluster host and prior to attempting any SunCluster configuration. For example:

```
Host "Sun1" with 2 HBAs: "qlc0" and "qlc1"
Host "Sun2" with 2 HBAs: "qlc0" and "qlc1"
DS8000 with 4 Target adapters:
  B1-C1 (cpsspc400)
  B2-C1 (cpsspc600)
  B3-C1 (cpsspc500)
  B4-C1 (cpsspc700)
```

An acceptable zoning configuration would be:

```
Sun1:qlc0 -- B1-C1
Sun1:qlc1 -- B4-C1
Sun2:qlc0 -- B2-C1
Sun2:qlc1 -- B3-C1
```

Note that this configuration uses only two target adapters, but each is paired with the HBAs in separate zones, which satisfies the requirement that each initiator be in a zone with itself and only 1 target.

Another acceptable zoning configuration would be:

```
zone 1: Sun1:qlc0 -- B1-C1
zone 2: Sun1:qlc1 -- B4-C1
zone 3: Sun2:qlc0 -- B2-C1
zone 4: Sun2:qlc1 -- B3-C1
```

This restriction is unique to *two-node* SunCluster 3.x. with MPxIO. The use of MPxIO is limited to Sun-branded host adapters, thus the restriction does not apply to other adapter brands (AMCC, Emulex, QLogic).

IBM will add more information as it becomes available. You can refer to the following references for more information:

- IBM DS8000 Interoperability Matrix, see <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>
- PMH 20238,500,624 (2105-800)
- PMH 20794,500,624 / 21374,500,624 (2042 CNT/Inrange Switch)
- CMT Complaint AQ4862

Chapter 19. Attaching to a Windows 2000 host

This chapter provides instructions, requirements and considerations for attaching a Windows 2000 host system to a storage unit.

Attaching with fibre-channel adapters

This section provides information for attaching a storage unit to a Windows 2000 host system with fibre-channel adapters. It also describes how to install, download and configure the adapter cards.

Supported fibre-channel adapters for Windows 2000

This section lists the supported fibre-channel adapters for the Windows 2000 host system.

The following adapters are supported for the Windows 2000 host System:

- Emulex LP9002L adapter card
- Emulex LP9002DC adapter card
- Emulex LP9402DC adapter card
- Emulex LP9802 adapter card
- Emulex LP10000 adapter card
- Emulex LP10000DC adapter card
- Netfinity adapter card
- QLogic QLA2310FL adapter card
- QLogic QLA2340FL adapter card
- QLogic QLA2342FL adapter card

The storage unit supports eServer BladeCenter HS20 with blades 8678, 8832, and 8839 in a Windows 2000, 32-bit environment.

The storage unit supports both direct and fabric attachment to the eServer BladeCenter JS20 when operating with AIX 5.2H, or Red Hat Enterprise Linux 3.0, or Linux SLES 8 or 9.

Note: For more information about BladeCenter HS20 and BladeCenter JS20 and for a list of open-systems hosts, operating systems, adapters and switches that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Fibre-channel attachment requirements for the Windows 2000 host system

This section lists the requirements for attaching the storage unit to your host system.

Review the following attachment requirements before attaching with fibre-channel adapters:

- Check the LUN limitations for your host system.
- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *User's Guide* is on the TotalStorage CD that you receive with the storage unit.

- See <http://www-1.ibm.com/servers/storage/disk/ds8000> for details about the release level for your operating system.

Either you or an IBM service support representative must perform the following tasks to install and configure a storage unit:

1. Install the storage unit.
2. Define the fibre-channel host system with the worldwide port name identifiers. For the list of worldwide port names, see “Locating the worldwide port name (WWPN),” on page 225.
3. Define the fibre-port configuration if you did not do it when you installed the storage unit or fibre-channel adapters.
4. Configure the host system for the storage unit by using the instructions in your host system publications.

Configuration Requirements

You must consider the following configuration requirements when you attach a storage unit to a Windows 2000 host system:

- Version 1.4.0.0 or later of the SDD supports the Windows Server 2000 host system in a clustering environment. To have failover protection on an open system, the SDD requires at least two paths, and SDD allows a maximum of 32 paths. See <http://www-1.ibm.com/servers/storage/support/software/sdd.html> for the most current information about SDD.
- Map to the LUNs for the target volumes of the Windows 2000 host when you need access to the data on the target volume. This greatly reduces the time it takes for the host system to restart. If you do not map the LUNs for the target volumes of the Windows 2000 host, the system can take up to 10 minutes per Remote Mirror and Copy target volume to map the LUNs.
- If you map the LUNs, perform the LUN mapping after the Remote Mirror and Copy operation and immediately before you access the data. You must restart the host system before you can access the data on the target volume.
- Perform the mapping when the target volume is mapped to the host through the IBM TotalStorage DS Storage Manager.

Fibre-channel attachment considerations for Windows 2000

This section lists the fibre-channel attachment considerations for your Windows 2000 host system.

The IBM TotalStorage DS8000 supports attachment to a Windows 2000 host running VERITAS Volume Manager 4.2 for clustered and nonclustered configurations.

For up-to-date information about the operating systems, hosts, adapter cards, switch, and fabric connections that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Installing and configuring the Emulex adapter card

The following sections describe how to download, install, and configure the Emulex adapter card for Windows 2000.

Installation considerations for the Emulex adapter card for Windows 2000

This section provides instructions for attaching a Windows 2000 host system to a storage unit with an Emulex adapter card.

Note: The Emulex LP9402DC adapter card is a 3.3 volt card only and requires a 3.3 volt PCI slot.

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

- Target
 - Public initiator
 - Private initiator
 - Target and public initiator
 - Target and private initiator
1. Record the IEEE number that is printed on the card. You can use the IEEE number to determine the WWPNN.
 2. Refer to the installation instructions provided by your host adapter vendor for their specific installation instructions. See <http://www.emulex.com> for the latest documentation.

Downloading the Emulex adapter driver for Windows 2000

You can use this information to download the Emulex adapter driver for Windows 2000.

Review the installation considerations for the Emulex adapter card before downloading the adapter driver.

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

1. Restart your host system.
2. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
3. Click **Interoperability matrix**.
4. Click **DS8000 interoperability matrix**.
5. Find the section for the current version of the driver and firmware and driver you want.
6. Go to <http://www.emulex.com>.
7. Click **Driver, downloads, and documentation** from the left navigation pane. Click **OEM software and documentation**.
8. Select **IBM** from the drop down vendor list.
9. Click the link for the adapter that you want to get the firmware, driver, and documentation to install and download the adapter driver.

Installing the Emulex adapter driver for Windows 2000

This section provides instructions for installing the Emulex adapter driver for Windows 2000.

Review the installation considerations for the Emulex adapter card before installing the adapter driver.

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

1. Restart your host system.
2. Go to www.emulex.com.
3. Click **drivers, software, and manuals** from the left navigation pane
4. In the **HBA Drivers** section, select **Windows** from the dropdown OS list.
5. Click **Windows 2000** for the FC Port Driver.
6. Click **Installation** to get the documentation that tells you how to install the adapter driver.

Configuring the Emulex adapter drivers for Windows 2000

You can use these recommended configuration settings for your Emulex adapters on a Windows 2000 host system.

Table 16 provides recommended configuration file parameters for an Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, or LP10000DC adapters. The settings are for a storage unit that is attached in one of the following ways:

- Through a switch using the fabric, automap, SCSI-device port driver
- Directly using the arbitrated loop automap, SCSI-device port driver

Note: For the most current information about fibre-channel adapter parameter settings, see

<http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do>

Table 16. Recommended configuration file parameters for the Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, and LP10000DC adapters

Parameters	Recommended settings
Automatically map SCSI devices	Checked (enabled)
Query name server for all N-ports	Checked (enabled)
Allow multiple paths to SCSI targets	Checked (enabled)
Register for state change	Checked (enabled)
Use report LUNs	Checked (enabled)
Use name server after RSCN	Checked (enabled)
LUN mapping	Checked (enabled)
Automatic LUN mapping	Checked (enabled)
Scan in device ID order	Not checked (disabled)
Disable target reset for tape devices	Not checked (disabled)
Use PLOGI instead of PDISC after LIP	Checked (enabled)
Enable class 2 for SCSI devices	Not checked (disabled)
Report unknown SCSI devices	Not checked (disabled)
Look for disappearing devices	Not checked (disabled)
Translate queue full to busy	Checked (enabled) if attached LUNs are part of a PPRC or FlashCopy pair. Not checked (disabled) if attached LUNs are not part of a PPRC or FlashCopy pair.

Table 16. Recommended configuration file parameters for the Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, and LP10000DC adapters (continued)

Parameters	Recommended settings
Use bus reset status for retries	Not checked (disabled)
Retry unit attention	Not checked (disabled)
Retry PLOGI open failures	Not checked (disabled)
Enable FCP-2 Recovery	Checked (enabled)
Maximum number of LUNs	Equal to or greater than the number of the storage unit LUNs available to the host bus adapter
Maximum queue depth	8
Link Timer	30 seconds
Retries	64
E_D_TOV	2000 milliseconds
AL_TOV	15 milliseconds
Wait ready timer	45 seconds
Retry timer	2000 milliseconds
R_A_TOV	2 seconds
ARB_TOV	1000 milliseconds
Link Control	
Topology	<ul style="list-style-type: none"> • Point-to-point (fabric) • Arbitrated loop (direct attachment)
Link speed	Auto
Host pointers in SLIM	Not checked (disabled)

Installing and configuring the Netfinity adapter card for Windows 2000

The following section describes how to download, install, and configure the Netfinity adapter card for Windows 2000.

Installing the Netfinity adapter card for Windows 2000

This section provides instructions for installing the Netfinity adapter card for the Windows 2000 host system.

1. Obtain an optical multimode cable with a standard-connector style of duplex connector.
2. Write down the serial number of the adapter card.
In the unlikely event that the nonvolatile random access memory is damaged, the system will prompt you for the Netfinityadapter serial number.
3. Check the system board and make any configuration changes necessary to accommodate the Netfinity adapter.
4. Turn off the peripheral devices and then the host.
5. Disconnect the power cord.
6. Remove the cover to the host.
7. Choose any PCI bus slot that supports 32- or 64-bit adapters.
You can install the adapter in a 64-bit PCI slot if one is available. If a 64-bit slot is not available, use a 32-bit PCI slot.

8. Align the adapter with the expansion slot, and carefully press the adapter into the slot until you hear the adapter snap into place.
9. Secure the adapter.
10. Connect one end of the fiber cable from the devices to the J1 connector on the adapter.
11. Carefully reinstall and secure the cover.
12. Connect the power cables.
13. Turn on all external fibre-channel devices.
14. Observe the monitor.

Note: You will not see the BIOS menu because the adapter BIOS is set to disabled by default. Instead, you should see the information in Figure 93.

```
QLogic Corporation
QLA2xxx PCI Fibre Channel ROM BIOS Version X.XX
Copyright (C) QLogic Corporation 1998 All rights reserved.
www.qlogic.com

Press <Alt+Q> for FAST!Util

BIOS for Adapter X is disabled.

ROM BIOS not installed
```

Figure 93. Example of what is displayed when start the Windows 2000 host

15. To see the devices, press Alt+Q, and then use the utility program.

Checking the level of the device driver for Netfinity adapter card for Windows 2000

This section provides instructions to check the version level of the device driver for your Netfinity adapter card.

1. Check www.ibm.com/pc/support for the latest versions of the device drivers, utilities, and documentation.
2. Create backup copies of the following driver diskettes that are included with the adapter card:
 - a. Netfinity host adapter for NetWare
 - b. Netfinity host adapter for Windows 2000
3. Read the software installation chapters to determine which device driver you want to install.
4. Place the device driver diskette in the diskette drive.
5. Copy the QLVER.EXE file from the diskette into the directory that contains the QLogic device drivers.
For example, type: `COPY A:\QLVER.EXE C:\QLC\QLVER.EXE`
6. Change directories to the subdirectory that contains the QLogic drivers.
For example, type: `CD \QLC`
7. At the command prompt, type: `QLVER *.*`
This action invokes the QLogic banner, followed by a list of drivers and version numbers.

Documentation that comes with your Netfinity adapter card

This section lists documentation that comes with your Netfinity adapter card.

The following files are on the diskette that is shipped with your adapter card:

Type	Description
TXTSETUP.OEM	The driver installation script for the initial Windows text setup
QLOGIC	The identification file for the Windows setup program.
\W2K\OEMSETUP.IN	The driver installation script for the Windows setup program
\W2K\QL2200.SYS	The Windows 2000 device driver
\W2K\README.TXT	Helpful hints about the Windows 2000 device driver
\W2K\RELEASE.TXT	The history of release fixes

Installing the Netfinity device driver for Windows 2000

This section provides instructions to install the Netfinity device driver on a host that already has Windows 2000 installed.

You must already have installed the Windows 2000 operating system installed on your host machine.

1. Install the Netfinity adapter.
2. Restart Windows 2000.
3. Insert the Netfinity host adapter device driver for Windows 2000 diskette and follow the instructions that are displayed.

Updating the Windows 2000 device driver

This section provides instructions to replace an existing Windows 2000 device driver with an updated version.

1. Right click **My Computer** → **Manage** → **Device Manager**.
2. Click the plus (+) sign to expand the menu for **SCSI and RAID Controllers**.
3. Right click on the adapter that you want to update and select **Properties**.
4. Select the tab for **Driver**.
5. Click **Update Driver** and follow the instructions that appear on the screen.

Installing and configuring the QLogic adapter cards in Windows 2000

This section describes how to install and configure the QLogic adapter cards for the Windows 2000 host system.

Installing the QLogic QLA23xx adapter card for Windows 2000

This section describes how to attach a storage unit to a Windows 2000 host system with the QLogic QLA23xx adapter card.

1. Install the QLogic QLA23xx adapter card in the host system.
2. Connect the cable to the storage unit port.
3. Restart the host.
4. Press Ctrl+Q to get to the **FAST!Util** menu.

5. Use FAST!Util to set the parameters to the values listed in the following table:

Table 17. Recommended settings for the QLogic QLA23xx adapter card for Windows 2000

Parameters	Recommended settings
Basic Input/Output System (BIOS)	Disabled
Frame size	2048
Loop reset delay	5 (minimum)
Adapter hard loop ID	Disabled
Hard loop ID	0
Execution throttle	100
Fast command posting	Enabled
>4 GB addressing	Disabled (for 32-bit systems)
LUNs per target	0
Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes Note: Enable LIP reset, Enable LIP full login, and Enable target reset parameters control the behavior of the adapter when Windows 2000 tries to do a SCSI bus reset. You must perform a target reset to make cluster failovers work. Use the SCSI bus device reset option to clear SCSI reservations.
Login retry count	20 (minimum)
Port down retry count	20 (minimum)
Driver load RISC code	Enabled
Enable database updates	No
Disable database load	No
IOCB allocation	256
Extended error logging	Disabled (You might want to enable this parameter if you use debugging.)

Table 17. Recommended settings for the QLogic QLA23xx adapter card for Windows 2000 (continued)

Parameters	Recommended settings
Type of connection	<ul style="list-style-type: none"> • Loop only • Point-to-point • Loop preferred (If you cannot use arbitrated loop, then default to point-to-point.) • Point-to point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop.) <p>Note:</p> <ol style="list-style-type: none"> 1. If you connect the storage unit directly to the host system, the option you select must match the port connections on the storage unit. 2. If you connect through a switch, the options do not need to match the port connections because the storage unit is point-to-point. 3. The appropriate host adapter on the server must also support a point-to-point connection on a direct connection. 4. If you use adapter cards from different manufacturers, they will not function properly in a direct point-to-point connection. This is not true if you connect through a switch because the storage unit is point-to-point.

6. Restart the server.

Downloading the QLogic adapter driver

This section provides instructions to load the current driver onto the QLogic adapter card.

1. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
2. Click **Interoperability matrix**.
3. Click **DS8000 interoperability matrix**.
4. Find the section for the current version of the driver and firmware and driver you want.
5. Go to <http://www.qlogic.com>
6. Click **Drivers/Software** from the **Quick Links** list.
7. Choose **IBM** from the **OEM and ISV approved Drivers and Firmware** section.
8. Navigate to the part of the page for the adapter that you want to download, then click **Download**.

Installing the QLogic adapter driver for Windows 2000

This section provides instructions to install the QLogic adapter drivers.

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

1. Go to <http://www.qlogic.com>

2. From the Quick Links dropdown list, click **Drivers/Software**.
3. In the HBA section, click the Down Arrow to highlight and click the fibre-channel adapter that you want and follow the instructions from the web site.

Verifying that Windows 2000 is configured for storage

This section provides instructions to determine if your Windows 2000 host system is configured for storage.

1. Partition new drives with Disk Management.
2. From the Windows 2000 desktop, right click **Start**.
3. Click **Explore** and verify that you can see the fibre-channel drives.
4. Select a large file (for example, 9 MB file), and drag (copy) it to a fibre-channel drive.

Configuring for availability and recoverability

This section provides availability and recoverability considerations for attaching a Windows 2000 host system to a storage unit.

Configuration considerations

This section describes how to ensure optimum availability and recoverability when you attach a storage unit to a Windows 2000 host system.

You must set the TimeOutValue value associated with the supported host adapters to 60 seconds. The setting is consistent with the configuration for FC-AL adapters and disk subsystems when attached to Windows 2000 host system.

The host adapter uses the TimeOutValue parameter to bind its recovery actions and responses to the disk subsystem. The value exists in different places in the system configuration. You can retrieve and use it in different ways depending on the type of host adapter that is installed. "Setting the TimeOutValue registry for Windows 2000" tells you how to modify the value safely in the Windows 2000 registry.

Setting the TimeOutValue registry for Windows 2000

This section provides instructions for setting the TimeOutValue registry for Windows 2000.

1. From the **Run** menu or command prompt, type:
Regedt32.exe
2. Navigate to the following registry key:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk
3. Look for the value called TimeOutValue. If the value called TimeOutValue does not exist, go to step 3a. If the TimeOutValue exists, go to step 4 on page 207.
 - a. Click **Edit → Add Value...**
 - b. For ValueName, type: TimeOutValue.
 - c. For data type, type: REG-DWORD.
 - d. Click **OK**.
 - e. For data, type: 3c.
 - f. For radix, click **Hex**.
 - g. Click **OK**.

4. If the value exists and is less than 0x0000003c (60 decimal), perform the following steps to increase it to 0x3c.
 - a. Click **TimeOutValue**.
 - b. Click **Edit → DWORD...**
 - c. For data, type: 3c.
 - d. For radix, click **Hex**.
 - e. Click **OK**.
5. Exit the Regedt32 program.
6. Restart your Windows 2000 server for the changes to take effect.

Installing remote fibre-channel boot support for a Windows 2000 host system

This section provides instructions for installing remote fibre-channel boot support for a Windows 2000 host system.

You must obtain the following items before beginning to install remote fibre-channel boot support:

- The *IBM TotalStorage Enterprise Storage Server Subsystem Device Driver User's Guide*.
- The *IBM TotalStorage DS8000 Introduction and Planning Guide*.
- Either the Windows 2000 CD-ROM or a set of Windows 2000 setup floppy disks with the Windows 2000 CD-ROM.
- The Windows 2000 installation instructions.
- Ensure that the host adapters are installed and cable connections of the host, switch, and storage unit are complete and functional.
- The supported fibre-channel host adapter driver package.
- The supported fibre-channel host adapter driver documentation.
- Extract and copy the driver files to a floppy.

Remote fibre-channel boot support is available for Windows 2000 with or without SDD. For information on the SDD hardware and software installation requirements, see *IBM TotalStorage Multipath Subsystem Device Driver User's Guide*.

Configure zoning and obtain storage

1. Reboot the host system to obtain the Host Bus Adapter WWPN:
 - a. **For QLogic host adapters:** A QLogic BIOS banner appears, press <Alt Q> or press <Ctrl Q> as indicated, depending on the BIOS level. For each host adapter, select Configuration Settings -> Host Bus Adapter Settings and write down the host adapter WWPN.
 - b. **For Emulex host adapters:** When the BIOS banner appears, press either <Alt E> or press <Alt F5> as indicated, depending on the BIOS level. A list of adapters is displayed. For each host adapter, select the adapter and write down the host adapter WWPN.
2. Configure appropriate zoning on the switch to include only the storage unit WWPN and one host adapter WWPN.
3. Use the IBM TotalStorage DS Storage Manager to add the host system to the list of host systems that are available for attachment to the storage unit. For detailed information of how to use the IBM TotalStorage DS Storage Manager, refer to the *IBM TotalStorage DS8000 Introduction and Planning Guide*.

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

Flash QLogic host adapter

Because the host adapter BIOS is responsible for scanning the fibre-channel bus to locate the remote boot device, check that the latest supported host adapter BIOS is loaded. If not, follow the host adapter documentation to update the BIOS of the host adapter.

Configure QLogic host adapters

1. Return to the host system.
2. **For QLogic host adapters:** When the QLogic BIOS banner appears, press <Ctrl Q>. The QLogic Fast!UTIL Menu appears. Depending on the host adapter BIOS version, the BIOS menu may be different in the following step:
 - a. Select Configuration Settings->Host Bus Adapter Settings. Set only one of the Host Adapters BIOS to **Enabled** and set **Connection Options** appropriately. For example, set to 1 for fabric point-to-point. Press ESC to return to the previous menu.
 - b. Select Configuration Settings->Selectable Boot Settings. Set Selectable Boot to **Enabled** and select the appropriate Boot Port Name and the LUN that corresponds to the WWPN of the storage unit.
3. **For Emulex host adapters:** When Emulex BIOS banner appears, press Alt E or press Alt F5 as indicated, depending on the BIOS level. A list of adapters is displayed. For each host adapter, select the adapter by entering the appropriate number.
 - a. Set only one host adapter to **Enabled** and set **Topology** appropriately. For example, set to 1 for fabric point-to-point. Press ESC to return to the previous menu.
 - b. Select Configure Boot Devices and configure one boot device, including the boot device, the boot entry WWPN, and the starting LUN.
4. Exit the BIOS utility and select to save changes.

Windows 2000 Installation

Apply the Windows Server 2000 installation instructions. These instructions include the following steps:

1. Reboot the host system and insert the Windows 2000 setup CD-ROM in the CD-ROM drive.
2. If prompted with the message, **Press any key to boot from CD**, press any key.
3. Respond to **Press F6 if you want to install a third party SCSI or RAID driver** message by pressing F6
4. On **Welcome to Setup** panel, press Enter to select the **Set up of Windows now** message.
5. Answer appropriately on the **Windows Licensing Agreement** panel.
6. Press ESC to reply to the **Continue installing a fresh copy of Windows** prompt.
7. When prompted, choose the disk on the storage unit where the Windows Server 2000 operating system will reside.

Windows 2000 Post Installation

1. Once Windows Server 2000 installation is complete, install the IBM Subsystem Device Driver (SDD). Refer to the *IBM TotalStorage Multipath Subsystem Device Driver User's Guide* and the SDD Readme.txt for the installation of SDD.
2. Update zoning on the switch to permit the availability of multiple paths.

Chapter 20. Attaching to a Windows Server 2003 host

This chapter provides instructions, requirements, and considerations for attaching a Windows Server 2003 host system to a storage unit.

Attaching with fibre-channel adapters

This section provides information for attaching a storage unit to a Windows Server 2003 host system with fibre-channel adapters. It also describes how to install, download, and configure the adapter cards.

Supported fibre-channel adapters for Windows Server 2003

This section lists the supported fibre-channel adapters for the Windows Server 2003 host system.

The following adapters are supported for the Windows Server 2003 host system:

- Emulex LP9002L adapter card
- Emulex LP9002DC adapter card
- Emulex LP9402DC adapter card
- Emulex LP9802 adapter card
- Emulex LP10000 adapter card
- Emulex LP10000DC adapter card
- Netfinity adapter card
- QLogic QLA2310FL adapter card
- QLogic QLA2340FL adapter card
- QLogic QLA2342FL adapter card

The storage unit supports eServer BladeCenter HS20 with blades 8678, 8832, and 8839 in a Windows Server 2003, 32-bit environment.

The storage unit supports both direct and fabric attachment to the eServer BladeCenter JS20 when operating with AIX 5.2H, or Red Hat Enterprise Linux 3.0, or Linux SLES 8 or 9.

Note: For more information about BladeCenter HS20 and BladeCenter JS20 and for a list of open-systems hosts, operating systems, adapters and switches that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Fibre-channel attachment requirements for the Windows Server 2003 host system

This section lists the requirements for attaching the storage unit to your host system.

Review the following attachment requirements before attaching with fibre-channel adapters:

- Check the LUN limitations for your host system.
- Ensure that you can reference the documentation for your host system and the *IBM TotalStorage DS8000 User's Guide*. The *User's Guide* is on the TotalStorage CD that you receive with the storage unit.

- See <http://www-1.ibm.com/servers/storage/disk/ds8000> for details about the release level for your operating system.

Either you or an IBM service support representative must perform the following tasks to install and configure a storage unit:

1. Install the storage unit.
2. Define the fibre-channel host system with the worldwide port name identifiers. For the list of worldwide port names, see “Locating the worldwide port name (WWPN),” on page 225.
3. Define the fibre-port configuration if you did not do it when you installed the storage unit or fibre-channel adapters.
4. Configure the host system for the storage unit by using the instructions in your host system publications.

Configuration requirements

You must consider the following configuration requirements when you attach a storage unit to a Windows Server 2003 host system:

- Version 1.4.0.0 or later of the SDD supports the Windows Server 2003 host system in a clustering environment. To have failover protection on an open system, SDD requires at least two paths, and SDD allows a maximum of 32 paths. See <http://www-1.ibm.com/servers/storage/support/software/sdd.html> for the most current information about SDD.
- Map to the LUNs for the target volumes of the Windows Server 2003 host when you need access to the data on the target volume. This greatly reduces the time it takes for the host system to restart. If you do not map the LUNs for the target volumes of the Windows Server 2003 host, the system can take up to 10 minutes per Remote Mirror and Copy target volume to map the LUNs.
- If you map the LUNs, perform the LUN mapping after the Remote Mirror and Copy operation and immediately before you access the data. You must restart the host system before you can access the data on the target volume.
- Perform the mapping when the target volume is mapped to the host through the IBM TotalStorage DS Storage Manager.

Fibre-channel attachment considerations for Windows Server 2003

This section lists the fibre-channel attachment considerations for your Windows Server 2003 host system.

The IBM TotalStorage DS8000 supports attachment to a Windows Server 2003 host running VERITAS Volume Manager 4.2 for clustered and nonclustered configurations.

For up-to-date information about the servers, operating systems, hosts, adapter cards, switches, and fabric connections that IBM supports, see the *Interoperability Matrix* at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

Installing and configuring the Emulex adapter card

The following sections describe how to download, install, and configure the Emulex adapter card for the Windows Server 2003.

Installation considerations for the Emulex adapter card for Windows Server 2003

This section lists considerations for attaching a Windows Server 2003 host system to a storage unit with an Emulex adapter card.

This section describes considerations for attaching a Windows Server 2003 host system to a storage unit with an Emulex adapter card.

Note: The Emulex LP9402DC adapter card is a 3.3 volt card only and requires a 3.3 volt PCI slot.

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

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

Record the IEEE number that is printed on the Emulex adapter card for later use during installation. You can use the IEEE number to determine the WWPN.

Refer to the installation instructions provided by your host adapter vendor for their specific installation instructions. See <http://www.emulex.com> for the latest documentation.

Downloading the Emulex adapter driver for Windows Server 2003

This section provides instructions for downloading the Emulex adapter driver for Windows Server 2003.

Review the installation considerations for the Emulex adapter card before downloading the adapter driver.

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

1. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
2. Click **Interoperability matrix**.
3. Click **DS8000 interoperability matrix**.
4. Find the section for the current version of the driver and firmware and driver that you want.
5. Go to <http://www.emulex.com>.
6. Click **Driver, downloads, and documentation** from the left navigation pane. Click **OEM software and documentation**.
7. Select **IBM** from the drop down vendor list.
8. Click the link for the adapter that you want to get the firmware, driver, and documentation to install and download the adapter driver.

Installing the Emulex adapter driver for Windows Server 2003

This section provides instructions for installing the Emulex adapter driver for Windows Server 2003.

Review the installation considerations for the Emulex adapter card before installing the adapter driver.

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

1. Go to:.

<http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do>

2. Click **drivers, software, and manuals** from the left navigation pane.
3. In the **HBA Drivers** section, select **Windows** from the drop down OS list.
4. Click **Windows Server 2003** for the FC Port Driver.
5. Click **Installation** to get the documentation that tells you how to install the adapter driver.

Configuring the Emulex adapter drivers for Windows Server 2003

You can use these recommended configuration settings for your Emulex adapters on a Windows 2003 host system.

Table 18 provides recommended configuration file parameters for an Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, or LP10000DC adapters. The settings relate to the storage unit that is attached in one of the following ways:

- Through a switch using the fabric, automap, SCSI-device port driver
- Directly using the arbitrated loop automap, SCSI-device port driver

Note: For the most current information about fibre-channel adapter parameter settings, see:

<http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do>

Table 18. Recommended configuration file parameters for the Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, and LP10000DC adapters

Parameters	Recommended Settings
Automatically map SCSI devices	Checked (enabled)
Query name server for all N-ports	Checked (enabled)
Allow multiple paths to SCSI targets	Checked (enabled)
Point-to-point	<ul style="list-style-type: none">• Not checked (disabled) for direct attach• Not shown for the fabric attach
Register for state change	Checked (enabled)
Use report LUNs	Checked (enabled)
Use name server after RSCN	Checked (enabled)
LUN mapping	Checked (enabled)
Automatic LUN mapping	Checked (enabled)
Scan in device ID order	Not checked (disabled)
Use PLOGI instead of PDISC after LIP	Checked (enabled)
Enable class 2 for SCSI devices	Not checked (disabled)
Report unknown SCSI devices	Not checked (disabled)
Look for disappearing devices	Not checked (disabled)
Translate queue full to busy	Checked (enabled)

Table 18. Recommended configuration file parameters for the Emulex LP9002L, LP9002DC, LP9402DC, LP9802, LP10000, and LP10000DC adapters (continued)

Parameters	Recommended Settings
Use bus reset status for retries	Not checked (disabled)
Retry unit attention	Not checked (disabled)
Retry PLOGI open failures	Not checked (disabled)
Enable FCP-2 Recovery	Checked (enabled)
Maximum number of LUNs	Equal to or greater than the number of the storage unit LUNs available to the host bus adapter
Maximum queue depth	8
Quickfail	1
Link Timer	30 seconds
Retries	64
E_D_TOV	2000 milliseconds
AL_TOV	15 milliseconds
Wait ready timer	45 seconds
Retry timer	2000 milliseconds
R_A_TOV	2 seconds
ARB_TOV	1000 milliseconds
Link Control	
Topology	<ul style="list-style-type: none"> • Point-to-point (fabric) • Arbitrated loop (direct attachment)
Link speed	Auto

Installing and configuring the Netfinity adapter card

The following section describes how to download, install, and configure the Netfinity adapter card for the Windows Server 2003.

Installing the Netfinity adapter card for Windows Server 2003

This section provides instructions for installing the Netfinity adapter card for the Windows Server 2003 host system.

1. Obtain an optical multimode cable with a standard connector-style of duplex connector.
2. Write down the serial number of the adapter card.
In the unlikely event that the nonvolatile random access memory is damaged, the system prompts you for the Netfinityadapter serial number.
3. Check the system board and make any configuration changes that are necessary to accommodate the Netfinityadapter.
4. Turn off the peripheral devices and then the host.
5. Disconnect the power cord.
6. Remove the cover to the host.
7. Choose any PCI bus slot that supports 32- or 64-bit adapters.
You can install the adapter in a 64-bit PCI slot, if one is available. If a 64-bit slot is not available, use a 32-bit PCI slot.

8. Align the adapter with the expansion slot, and carefully press the adapter into the slot until you hear the adapter snap into place.
9. Secure the adapter.
10. Connect one end of the fiber cable from the devices to the J1 connector on the adapter.
11. Carefully reinstall and secure the cover.
12. Connect the power cables.
13. Turn on all external fibre-channel devices.
14. Observe the monitor.

Note: You will not see the BIOS menu because the adapter BIOS is set to disabled by default. Instead, you should see the information in Figure 94.

```
QLogic Corporation
QLA2xxx PCI Fibre Channel ROM BIOS Version X.XX
Copyright (C) QLogic Corporation 1998 All rights reserved.
www.qlogic.com

Press <Alt+Q> for FAST!Util

BIOS for Adapter X is disabled.

ROM BIOS not installed
```

Figure 94. Example of what the system displays when you start the Windows Server 2003 host

15. To see the devices, press Alt+Q, and then use the utility program.

Checking the level of the device driver for the Netfinity adapter card for Windows Server 2003

This section provides instructions to check the version of the device driver that is installed on your host.

1. Check www.ibm.com/pc/support for the latest versions of the device drivers, utilities, and documentation.
2. Create backup copies of the following driver diskettes that are included with the adapter card:
 - a. Netfinity host adapter for NetWare
 - b. Netfinity host adapter for Windows Server 2003
3. Read the software installation chapters to determine which device driver you want to install.
4. Place the device driver diskette in the diskette drive.
5. Copy the QLVER.EXE file from the diskette into the directory that contains the QLogic device drivers.
For example, type: `COPY A:\QLVER.EXE C:\QLC\QLVER.EXE`
6. Change directories to the subdirectory that contains the QLogic drivers.
For example, type: `CD \QLC`
7. At the command prompt, type: `QLVER *.*`
This action invokes the QLogic banner, followed by a list of drivers and version numbers.

Documentation that comes with your Netfinity adapter card

This section lists documentation that comes with your Netfinity adapter card.

The diskette that is shipped with your adapter card contains the following files:

Type	Description
TXTSETUP.OEM	The driver installation script for the initial Windows text setup
QLOGIC	The identification file for the Windows setup program
\W2K\OEMSETUP.INF	The driver installation script for the Windows setup program
\W2K\QL2200.SYS	The Windows Server 2003 device driver
\W2K\README.TXT	Helpful hints about the Windows Server 2003 device driver
\W2K\RELEASE.TXT	The history of release fixes

Installing the Netfinity device driver for Windows Server 2003

This section provides instructions to install the Netfinity device driver on a host that already has Windows Server 2003 installed.

You must already have installed the Windows Server 2003 operating system on your host machine.

These steps allow you to install the device driver on a host that already has installed Windows Server 2003.

1. Install the Netfinity adapter.
2. Restart Windows Server 2003.
3. Insert the Netfinity host adapter device driver for Windows Server 2003 diskette and follow the instructions that are displayed.

Updating the Windows Server 2003 device driver

This section provides instructions to replace an existing Windows Server 2003 device driver with an updated version.

1. Right click **My Computer** → **Manage** → **Device Manager**.
2. Click the plus (+) sign to expand the menu for **SCSI and RAID Controllers**.
3. Right click on the adapter that you want to update. Select **Update Driver**.
4. Follow the instructions on the screen.

Installing and configuring the QLogic adapter cards

This section describes how to install and configure the QLogic adapter cards for the Windows Server 2003 host system.

Installing the QLogic QLA23xx adapter card for Windows Server 2003

This section describes how to attach a storage unit to a Windows Server 2003 host system with the QLogic QLA23xx adapter card.

1. Install the QLogic QLA23xx adapter card in the host system.
2. Connect the cable to the storage unit port.

3. Restart the host.
4. Press Ctrl+Q to get to the **FAST!Util** menu.
5. Use FAST!Util to set the parameters to the values listed in the following table:

Table 19. Recommended settings for the QLogic QLA23xx adapter card for a Windows Server 2003 host

Parameters	Recommended settings
Basic Input/Output System (BIOS)	Disabled
Frame size	2048
Loop reset delay	5 (minimum)
Adapter hard loop ID	Disabled
Execution throttle	100
Fast command posting	Enabled
>4 GB addressing	Disabled (for 32-bit systems)
LUNs per target	0
Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes Note: Enable LIP reset, Enable LIP full login, and Enable target reset parameters control the behavior of the adapter when Windows Server 2003 tries to do a SCSI bus reset. You must perform a target reset to make cluster failovers work. Use the SCSI bus device reset option to clear SCSI reservations.
Login retry count	20 (minimum)
Port down retry count	20 (minimum)
Driver load RISC code	Enabled
Enable database updates	No
Disable database load	No
IOCB allocation	256
Extended error logging	Disabled (You might want to enable this parameter if you use debugging.)

Table 19. Recommended settings for the QLogic QLA23xx adapter card for a Windows Server 2003 host (continued)

Parameters	Recommended settings
Type of connection	<ul style="list-style-type: none"> • Loop only • Point-to-point • Loop preferred (If you cannot use arbitrated loop, then default to point-to-point.) • Point-to point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop.) <p>Note:</p> <ol style="list-style-type: none"> 1. If you connect the storage unit directly to the host system, the option that you select must match the port connections on the storage unit. 2. If you connect through a switch, the options do not need to match the port connections because the storage unit is point-to-point. 3. The appropriate host adapter on the server must also support a point-to-point connection on a direct connection. 4. If you use adapter cards from different manufacturers, they will not function properly in a direct point-to-point connection. This is not true if you connect through a switch because the storage unit is point-to-point.

6. Restart the server.

Downloading the QLogic adapter driver

This section provides instructions to load the current driver onto the QLogic adapter card.

1. Go to <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.
2. Click **Interoperability matrix**.
3. Click **DS8000 interoperability matrix**.
4. Find the section for the current version of the driver and firmware and driver that you want.
5. Go to <http://www.qlogic.com>
6. Click **Drivers / Software** from the **Quick Links** list.
7. Choose **IBM** from the **OEM and ISV approved Drivers and Firmware** section.
8. Navigate to the part of the page for the adapter that you want to download, and then click **Download**.

Installing the QLogic adapter driver for Windows Server 2003

This section provides instructions to install the QLogic adapter drivers.

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

1. Go to <http://www.qlogic.com>

2. From the Quick Links dropdown list, choose **Drivers / Software**.
3. In the **HBA** section, highlight and click the fibre-channel adapter that you want and follow the instructions from the web site.

Verifying that Windows Server 2003 is configured for storage

This section provides instructions to determine if your Windows Server 2003 host system is configured for storage.

1. Partition new drives with the Disk Management utility.
2. From the Windows Server 2003 desktop, right click **Start**.
3. Click **Explore** and verify that you can see a list of the fibre-channel drives.
4. Select a large file (for example, 9 MB file), and drag (copy) it to a fibre-channel drive.

Configuring for availability and recoverability

This section provides availability and recoverability considerations for attaching a Windows Server 2003 host system to a storage unit.

Configuration considerations

This section describes configuration considerations when attaching a storage unit to a Windows Server 2003 host system.

This section describes how to ensure the optimum ability to retrieve and recover data when you attach a Windows Server 2003 host system to a storage unit. You must set the timeout value that is associated with the supported host bus adapters to 60 seconds. The setting is consistent with the configuration for FC-AL adapters and disk subsystems when they are attached to Windows Server 2003 host system.

The host bus adapter uses the timeout parameter to bind its recovery actions and responses to the disk subsystem. The value exists in different places in the system configuration. You can retrieve and use it in different ways depending on the type of host bus adapter. See "Setting the TimeOutValue registry for Windows Server 2003" for instructions about how to modify the value safely in the Windows Server 2003 registry.

Setting the TimeOutValue registry for Windows Server 2003

This section provides instructions for setting the TimeOutValue registry for Windows Server 2003.

1. From the **Run** menu or command prompt, type:
Regedt32.exe
2. Navigate to the following registry key:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk
3. Look for the value called TimeOutValue. If the value called TimeOutValue does not exist, go to step 3a. If the value called TimeOutValue exists, go to step 4 on page 221.
 - a. Click **Edit → Add Value....**
 - b. For ValueName, type: TimeOutValue.
 - c. For data, type: REG_DWORD.
 - d. Click **OK**.
 - e. For data, type: 3c.

- f. For radix, click **Hex**.
 - g. Click **OK**.
4. If the **TimeOutValue** value exists and is less than 0x0000003c (60 decimal), perform the following steps to increase it to 0x3c.
 - a. Click **TimeOutValue**.
 - b. Click **Edit → DWORD...**
 - c. For data, type 3c.
 - d. For radix, click **hex**.
 - e. Click **OK**.
5. Exit the Regedt32 program.
6. Restart your Windows Server 2003 server for the changes to take effect.

Installing remote fibre-channel boot support for a Windows Server 2003 host system

This section provides instructions for installing remote fibre-channel boot support for a Windows Server 2003 host system.

You must obtain the following items before you can install remote fibre-channel boot support:

- The *IBM TotalStorage Enterprise Storage Server Subsystem Device Driver User's Guide*.
- The *IBM TotalStorage DS8000 Introduction and Planning Guide*.
- Either the Windows 2003 CD-ROM or a set of Windows 2003 setup floppy disks with the Windows 2003 CD-ROM.
- The Windows 2003 installation instructions.
- Ensure that the host adapters are installed and cable connections of the host, switch, and storage unit are complete and functional.
- The supported fibre-channel host adapter driver package.
- The supported fibre-channel host adapter driver documentation.
- Extract and copy the driver files to a floppy.

Remote fibre-channel boot support is available for Windows 2003 with or without SDD. For information on the SDD hardware and software installation requirements, see *IBM TotalStorage Multipath Subsystem Device Driver User's Guide*.

Configure zoning and obtain storage

1. Reboot the host system to obtain the Host Bus Adapter WWPN:
 - a. **For QLogic host adapters:** A QLogic BIOS banner appears, press <Alt Q> or press <Ctrl Q> as indicated, depending on the BIOS level. For each host adapter, select Configuration Settings -> Host Bus Adapter Settings and write down the host adapter WWPN.
 - b. **For Emulex host adapters:** When the BIOS banner appears, press either <Alt E> or press <Alt F5> as indicated, depending on the BIOS level. A list of adapters is displayed. For each host adapter, select the adapter and write down the host adapter WWPN.
2. Configure appropriate zoning on the switch to include only the storage unit WWPN and one host adapter WWPN.
3. Use the IBM TotalStorage DS Storage Manager to add the host system to the list of host systems that are available for attachment to the storage unit. For

detailed information of how to use the IBM TotalStorage DS Storage Manager, refer to the *IBM TotalStorage DS8000 Introduction and Planning Guide*.

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

Flash QLogic host adapter

Because the host adapter BIOS is responsible for scanning the fibre-channel bus to locate the remote boot device, check that the latest supported host adapter BIOS is loaded. If not, follow the host adapter documentation to update the BIOS of the host adapter.

Configure QLogic host adapters

1. Return to the host system.
2. **For QLogic host adapters:** When the QLogic BIOS banner appears, press <Ctrl Q>. The QLogic Fast!UTIL Menu appears. Depending on the host adapter BIOS version, the BIOS menu may be different in the following step:
 - a. Select Configuration Settings->Host Bus Adapter Settings. Set only one of the Host Adapters BIOS to **Enabled** and set **Connection Options** appropriately. For example, set to 1 for fabric point-to-point. Press ESC to return to the previous menu.
 - b. Select Configuration Settings->Selectable Boot Settings. Set Selectable Boot to **Enabled** and select the appropriate Boot Port Name and the LUN that corresponds to the WWPN of the storage unit.
3. **For Emulex host adapters:** When Emulex BIOS banner appears, press Alt E or press Alt F5 as indicated, depending on the BIOS level. A list of adapters is displayed. For each host adapter, select the adapter by entering the appropriate number.
 - a. Set only one host adapter to **Enabled** and set **Topology** appropriately. For example, set to 1 for fabric point-to-point. Press ESC to return to the previous menu.
 - b. Select Configure Boot Devices and configure one boot device, including the boot device, the boot entry WWPN, and the starting LUN.
4. Exit the BIOS utility and select to save changes.

Windows 2003 Installation

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

1. Reboot the host system and insert the Windows 2003 setup CD-ROM in the CD-ROM drive.
2. If prompted with the message, **Press any key to boot from CD**, press any key.
3. Respond to **Press F6 if you want to install a third party SCSI or RAID driver** message by pressing F6
4. On **Welcome to Setup** panel, press Enter to select the **Set up of Windows now** message.
5. Answer appropriately on the **Windows Licensing Agreement** panel.
6. Press ESC to reply to the **Continue installing a fresh copy of Windows** prompt.
7. When prompted, choose the disk on the storage unit where the Windows Server 2003 operating system will reside.

Windows 2003 Post Installation

1. Once Windows Server 2003 installation is complete, install the IBM Subsystem Device Driver (SDD). Refer to the *IBM TotalStorage Multipath Subsystem Device Driver User's Guide* and the SDD Readme.txt for the installation of SDD.
2. Update zoning on the switch to permit the availability of multiple paths.

Appendix. Locating the worldwide port name (WWPN)

This chapter tells you how to locate the WWPN value for a fibre-channel adapter on the following host systems:

- Fujitsu PRIMEPOWER
- Hewlett-Packard AlphaServer
- Hewlett-Packard 9000
- IBM eServer iSeries
- IBM eServer pSeries or RS/6000
- Linux
- Novell NetWare
- SGI (Silicon Graphics, Inc.)
- Sun
- Windows 2000
- Windows 2003

Fibre-channel port name identification

This section describes how to identify a fibre-channel port name when you are locating the WWPN.

The worldwide port name (WWPN) consists of exactly 16 hexadecimal characters (0 - 9 and A - F). The storage unit uses it to uniquely identify the fibre-channel adapter card that is installed in your host system. The storage unit automatically finds the WWPN for your host fibre-channel adapter when you attach your host system to the storage unit.

Note: If your host uses more than one fibre-channel adapter to connect to your storage unit, you must add multiple entries to the host list for this host. You must add one for each fibre-channel adapter. Each adapter will have its own unique WWPN.

The format and content of the fibre-channel port identifier are determined by the manufacturer of the link control facility for the applicable fibre-channel port. The identifier is an eight-byte field, which the fibre-channel protocols use to uniquely identify the fibre-channel port.

You can manually locate a unique worldwide port name for the storage unit by performing the steps in the following sections.

Locating the WWPN for a Fujitsu PRIMEPOWER host

This section lists the steps you must take to locate the WWPN for a PRIMEPOWER host system.

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

Locating the WWPN for a Hewlett-Packard AlphaServer host

This section provides instructions to locate the WWPN for a Hewlett-Packard AlphaServer host.

1. From the console prompt, type: `P0>>>wwidmgr -show ada`
See Figure 95 for an example of what you see when you type the **P0>>>wwidmgr -show** command.

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

Figure 95. Example of the output from the Hewlett-Packard AlphaServer `wwidmgr -show` command

You might receive the following errors:

Message:

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

Action:

Type `P00>>>init wwidmgr` again.

Message:

`wwidmgr`: No such command

Action:

Type `P00>>>set mode diag wwidmgr`

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

2. Type: `#fgrep wwn /var/adm/messages`

Figure 96 shows an example of the output when you type `#fgrep wwn /var/adm/messages`. You can find the WWPN in the last column.

```
...
Nov  9 09:01:16 osp1cpq-ds20 vmunix: KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
Nov 10 10:07:12 osp1cpq-ds20 vmunix: KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
Nov 13 17:25:28 osp1cpq-ds20 vmunix: KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
Nov 14 11:08:16 osp1cpq-ds20 vmunix: KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
Nov 15 10:49:31 osp1cpq-ds20 vmunix: KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
...
```

Figure 96. Example of the output from the Hewlett-Packard `#fgrep wwn /var/adm/messages` command

Locating the WWPN for a Hewlett-Packard host

This section provides instructions to locate the WWPN for a Hewlett-Packard host.

1. Go to the root directory.
2. Type: `ioscan -fn | more`
3. Look under the description for the Fibre Channel Mass Storage adapter.
For example, look for the device path name `/dev/td1` or `/dev/fcms1`.
4. Type: `fcmsutil /dev/td1 | grep world` where `/dev/td1` is the path.

Figure 97 shows an example of the output when you type `ioscan -fn | more`.

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

Figure 97. Example of the output from the Hewlett-Packard `fcmsutil /dev/td1 | grep world` command.

Locating the WWPN for an IBM eServer iSeries host

This topic provides instructions to locate the WWPN for an IBM iSeries host.

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

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

Note: If you have exchanged a 2766 fibre-channel IOA in the iSeries host or the IBM 2107 storage unit, you must update the WWPN of the new 2766 IOA.

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

This topic provides instructions to locate the WWPN for an IBM-pSeries or an RS/6000 host.

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

The network address is the fibre-channel adapter port WWPN value.

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

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

Locating the WWPN for a Linux host

This topic provides instructions to locate the worldwide port name (WWPN) for an Intel server running either the Red Hat or SUSE version of the Linux operating system with a QLogic adapter.

1. To locate the WWPN without restarting the host system You can find the WWPN in your `/proc/scsi/qlaadapter_type/n` directory, where *adapter_type* is the host adapter type and *n* is the host adapter number for your card. Your host adapter type can be either `qlaxxxx` for QLogic adapters or `lpfc` for Emulex adapters.
2. To locate the WWPN with restarting the host system
 - a. Restart the server.
 - b. Press Alt+Q to get the **FAST!Util** menu.

Note: If you have more than one fibre-channel adapter installed, all the fibre-channel adapters display. Scroll down to the adapter you want. Press Enter.

- c. From the **FAST!Util** menu, scroll down and select **Select Host Adapter**.
- d. Scroll up and highlight **Configuration Settings**. Press Enter.
- e. From the **Configuration Settings** menu, click **Host Adapter Settings**.
- f. Write down the 16-digit alphanumeric string that is displayed on your screen.

When you configure storage on the storage unit for fibre-channel attachment, you must specify the worldwide name (WWN) of the fibre-channel adapter in the IBM TotalStorage DS Storage Manager.

On Linux systems, the required WWN displays in the `/var/log/messages` system log file when you load the driver. This location, however, depends on the system configuration. You can perform this task automatically during startup, or you can manually perform the task to load the module using the **insmod** or **modprobe** command.

Figure 98 on page 229 shows an example of what is displayed when you type `more /var/log/messages`.

```
# more /var/log/messages
...
Nov 2 08:55:05 skywalker kernel: (scsi): Found a QLA2200 @ bus 1, device 0x1,
irq 20, iobase 0x2300
Nov 2 08:55:10 skywalker kernel: scsi(3): Configure NVRAM parameters...
Nov 2 08:55:10 skywalker kernel: scsi(3): Verifying loaded RISC code...
Nov 2 08:55:10 skywalker kernel: scsi(3): Verifying chip...
Nov 2 08:55:10 skywalker kernel: scsi(3): Waiting for LIP to complete...
Nov 2 08:55:10 skywalker kernel: scsi(3): LOOP UP detected
Nov 2 08:55:10 skywalker kernel: scsi3: Topology - (F_Port), Host Loop address 0xffff
Nov 2 08:55:10 skywalker kernel: scsi(3): Waiting for LIP to complete...
Nov 2 08:55:10 skywalker kernel: scsi3: Topology - (F_Port), Host Loop address 0xffff
Nov 2 08:55:10 skywalker kernel: scsi-qla1-adapter-node=200000e08b02e534;
Nov 2 08:55:10 skywalker kernel: scsi-qla1-adapter-port=210000e08b02e534;
Nov 2 08:55:10 skywalker kernel: scsi-qla1-target-0=5005076300c08b1f;
...
```

Figure 98. Example of what is displayed in the `/var/log/messages` file

Locating the WWPN for a Novell NetWare host

This section provides instructions to locate the WWPN for a Novell NetWare host.

1. Restart the server.
2. Press Alt+Q to get the **FAST!Util** menu.

Note: If you have more than one fibre-channel adapter installed, all the adapters display on the screen. Scroll down to the adapter you want. Press Enter.

3. From the **FAST!Util** menu, scroll down and select **Select Host Adapter**.
4. Scroll up and highlight **Configuration Settings**. Press Enter.
5. From the **Configuration Settings** menu, click **Host Adapter Settings**.
6. Write down the 16-digit alphanumeric string that is displayed on your screen.

Locating the WWPN for an SGI host

This section provides instructions to locate the WWPN for an SGI host.

1. Restart the server.
2. Type the **scsiha -w [bus_number | device]** command.

For example, type `scsiha -w 6 7 8 9`. Figure 99 shows an example of what you see.

```
# scsiha -w 6 7 8 9
6 Portname: 210000e08b05d207
7 Portname: 210000e08b04d539
8 Portname: 210000e08b050808
9 Portname: 210000e08b038fe6
#
```

Figure 99. Example of the `scsiha — bus_number device | command`

Locating the WWPN for a Sun host

This section provides instructions to locate the WWPN for a Sun host.

Note: If you have multiple host adapters installed, you will see more than one WWPN.

Perform the following steps to locate the worldwide port name (WWPN) for the following adapters:

- AMCC PCI adapter
 - AMCC SBUS adapter
 - QLogic adapter
 - Emulex adapter
1. After you install the adapter and you restart the host system, view the `/var/adm/messages` file.
 2. Search for the line that contains the following phrase:
 - a. For the AMCC SBUS adapter, search for `fcawx: Fibre Channel WWNN`, where `x` is the adapter number (0, 1, and so on). You can find the WWPN on the same line immediately after the WWNN.
 - b. For the AMCC PCI adapter, search for `fca-pcix: Fibre Channel WWNN`, where `x` is the adapter number (0, 1, and so on). You can find the WWPN on the same line following the WWNN.
 - c. For the 2Gb AMCC adapter, search for `AMCCc146xN: Fibre Channel WWNN`, where `N` is the adapter number (0, 1, and so on). You can find the WWPN on the same line following the WWNN.
 - d. For the QLogic QLA23xx adapter, search for `qla23200-hbax-adapter-port-name` where `x` is the adapter number (0, 1, and so on).
 - e. For the Emulex adapter, search for `lpfcx: Fibre Channel WWNN`, where `x` is the adapter number 0, 1, and so on).

Locating the WWPN for a Windows 2000 host

This section provides instructions to locate the worldwide port name (WWPN) for a Windows 2000 host.

This section provides instructions for locating the WWPN for Windows 2000 hosts with either the QLogic or the Emulex adapter.

1. Perform the following steps to locate the WWPN for a Windows 2000 host system with a QLogic adapter:
 - a. Restart the server.
 - b. Press Alt+Q to get the **FAST!Util** menu.

Note: If you have more than one fibre-channel adapter installed, all the fibre-channel adapters display. Scroll down to the adapter you want. Press Enter.
 - c. From the **FAST!Util** menu, scroll down and select **Select Host Adapter**.
 - d. Scroll up and highlight **Configuration Settings**. Press Enter.
 - e. From the **Configuration Settings** menu, click **Host Adapter Settings**.
 - f. Write down the 16-digit alphanumeric string that is displayed on your screen.
2. Perform the following steps to locate the WWPN for a Windows 2000 host system with an Emulex adapter:

- a. Click **Start → Programs → Emulex Configuration Tool**.
- b. From the Emulex Configuration Tool menu in the Available Adapters window, double-click the adapter entry for which you want to display the WWPN information.

Locating the WWPN for a Windows Server 2003 host

This section provides instructions to locate the worldwide port name (WWPN) for a Windows Server 2003 host.

This section provides instructions for locating the WWPN for Windows Server 2003 hosts with either the QLogic or the Emulex adapter.

1. Perform the following steps to locate the WWPN for a Windows Server 2003 host system with a QLogic adapter:
 - a. Restart the server.
 - b. Press Alt+Q to get the **FAST!Util** menu.

Note: If you have more than one fibre-channel adapter installed, all the fibre-channel adapters display. Scroll down to the adapter you want. Press Enter.
 - c. From the **FAST!Util** menu, scroll down and select **Select Host Adapter**.
 - d. Scroll up and highlight **Configuration Settings**. Press Enter.
 - e. From the **Configuration Settings** menu, click **Host Adapter Settings**.
 - f. Write down the 16-digit alphanumeric string that is displayed on your screen.
2. Perform the following steps to locate the WWPN for a Windows Server 2003 host system with an Emulex adapter:
 - a. Click **Start → Programs → Emulex Configuration Tool**.
 - b. From the Emulex Configuration Tool menu in the Available Adapters window, double-click the adapter entry for which you want to display the WWPN information.

Accessibility

Accessibility features provide users who have disabilities with the ability to successfully access information and use technology.

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features

These are the major accessibility features in the IBM TotalStorage DS8000 information:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. IBM Home Page Reader version 3.0 has been tested.
- You can operate features using the keyboard instead of the mouse.

Navigating by keyboard

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the IBM TotalStorage DS8000 information from the keyboard by using the shortcut keys for your browser or Home Page Reader. See your browser Help for a list of shortcut keys that it supports. See the following Web site for a list of shortcut keys supported by Home Page Reader: http://www-306.ibm.com/able/solution_offerings/keyshort.html

Accessing the publications

You can find HTML versions of the IBM TotalStorage DS8000 information at the following Web site:
<http://www.ehone.ibm.com/public/applications/publications/cgibin/pbi.cgi>

You can access the information using IBM Home Page Reader 3.0.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Terms and conditions for downloading and printing publications

Permissions for the use of the information you have selected for download are granted subject to the following terms and conditions and your indication of acceptance thereof.

Personal Use: You may reproduce this information for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of this information, or any portion thereof, without the express consent of the manufacturer.

Commercial Use: You may reproduce, distribute and display this information solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of this information, or reproduce, distribute or display this information or any portion thereof outside your enterprise, without the express consent of the manufacturer.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the information or any data, software or other intellectual property contained therein.

The manufacturer reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the information is detrimental to its interest or, as determined by the manufacturer, the above instructions are not properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations. THE MANUFACTURER MAKES NO GUARANTEE ABOUT THE CONTENT OF THIS INFORMATION. THE INFORMATION IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

All material copyrighted by IBM Corporation.

By downloading or printing information from this site, you have indicated your agreement with these terms and conditions.

Trademarks

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

- AIX
- DB2
- DFSMS/MVS
- DFSMS/VM
- e (logo)
- Enterprise Storage Server
- ES/9000
- ESCON
- FICON
- FlashCopy
- HACMP
- IBM
- Intellistation
- MVS/ESA
- Netfinity
- NetVista
- Operating System/400
- OS/400
- RS/6000
- S/390
- Seascape
- SNAP/SHOT
- SP
- System/390
- TotalStorage
- Versatile Storage Server
- Virtualization Engine
- VSE/ESA
- z/Architecture
- z/OS
- z/VM
- zSeries

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

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

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission (FCC) statement

This equipment has been tested and complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the users authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conform à la norme NMB-003 du Canada.

European community compliance statement

This product is in conformity with the protection requirements of EC Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

Germany only

Zulassungsbescheinigung laut Gesetz ueber die elektromagnetische Vertraeglichkeit von Geraeten (EMVG) vom 30. August 1995.

Dieses Geraet ist berechtigt, in Uebereinstimmung mit dem deutschen EMVG das EG-Konformitaetszeichen - CE - zu fuehren.

Der Aussteller der Konformitaetserklaeung ist die IBM Deutschland.

Informationen in Hinsicht EMVG Paragraph 3 Abs. (2) 2:

Das Geraet erfuehlt die Schutzanforderungen nach EN 50082-1 und EN 55022 Klasse A.

EN 55022 Klasse A Geraete beduerfen folgender Hinweise:

Nach dem EMVG:

"Geraete duerfen an Orten, fuer die sie nicht ausreichend entstoert sind, nur mit besonderer Genehmigung des Bundesministeriums fuer Post und Telekommunikation oder des Bundesamtes fuer Post und Telekommunikation betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Stoerungen zu erwarten sind." (Auszug aus dem EMVG, Paragraph 3, Abs.4)

Dieses Genehmigungsverfahren ist nach Paragraph 9 EMVG in Verbindung mit der entsprechenden Kostenverordnung (Amtsblatt 14/93) kostenpflichtig.

Nach der EN 55022:

"Dies ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funkstoerungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Massnahmen durchzufuehren und dafuer aufzukommen."

Anmerkung:

Um die Einhaltung des EMVG sicherzustellen, sind die Geraete wie in den Handbuechern angegeben zu installieren und zu betreiben.

Japanese Voluntary Control Council for Interference (VCCI) class A statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean Ministry of Information and Communication (MIC) statement

Please note that this device has been certified for business use with regard to electromagnetic interference. If you find this is not suitable for your use, you may exchange it for one of residential use.

Taiwan class A compliance statement

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

VS07171L

Glossary

This glossary includes terms for the IBM TotalStorage and other Resiliency Family products.

This glossary includes selected terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172–1990, copyright 1990 by the American National Standards Institute (ANSI), 11 West 42nd Street, New York, New York 10036. Definitions derived from this book have the symbol (A) after the definition.
- *IBM Terminology*, which is available online at the following Web site: <http://www-306.ibm.com/ibm/terminology/index.html>. Definitions derived from this source have the symbol (GC) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions derived from this book have the symbol (I) after the definition. Definitions taken from draft international standards, committee drafts, and working papers that the ISO/IEC JTC1/SC1 is developing have the symbol (T) after the definition, indicating that final agreement has not been reached among the participating National Bodies of SC1.

This glossary uses the following cross-reference forms:

- See** Refers the reader to one of two kinds of related information:
- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
 - A synonym or more preferred term

See also Refers the reader to one or more related terms.

Contrast with

Refers the reader to a term that has an opposite or substantively different meaning.

Numerics

- 750** A model of the Enterprise Storage Server featuring a 2-way processor with limited physical storage capacity. This model can be updated to the model 800.
- 800** A model of the Enterprise Storage Server featuring a standard processor or an optional Turbo processor. The Model 800 supports RAID 5, RAID 10, and 15000 rpm drives. Model 800 supersedes Model F20.
- 1750** The machine type for the IBM TotalStorage DS6000 series. Models for the DS6000 include the 511 and EX1.
- 2105** The machine number for the IBM TotalStorage Enterprise Storage Server. Models of the Enterprise Storage Server are expressed as the number 2105 followed by "Model <xxx>", such as 2105 Model 800. The 2105 Model 100 is an Enterprise Storage Server expansion enclosure that is typically referred to simply as the Model 100.
- 2107** The machine type for the IBM TotalStorage DS8000 series. Models for the DS8000 series include the base units 921, 922 and 9A2 and expansion units 92E and 9AE.
- 3390** The machine number of an IBM disk storage system. The Enterprise Storage Server, when interfaced to IBM zSeries hosts, is set up to appear as one or more 3390 devices, with a choice of 3390-2, 3390-3, or 3390-9 track formats.
- 3990** The machine number of an IBM control unit.
- 7133** The machine number of an IBM disk storage system. The Model D40 and 020 drawers of the 7133 can be installed in the 2105-100 expansion enclosure of the ESS.

A

access

- 1) To obtain the use of a computer resource.
- 2) In computer security, a specific type of interaction between a subject and an object that results in flow of information from one to the other.

access-any mode

One of the two access modes that can be set for the storage unit during initial configuration. It enables all fibre-channel-attached host systems with no defined access profile to access all logical volumes on the storage unit. With a profile defined in DS Storage Manager for a particular host, that host has access only to volumes that are assigned to the WWPN for that host. See also *pseudo host* and *worldwide port name*.

ACK See *request for acknowledgment and acknowledgment*.

agent A program that automatically performs some service without user intervention or on a regular schedule. See also *subagent*.

alert A message or log that a storage unit generates as the result of error event collection and analysis. An alert indicates that a service action is required.

allegiance

For zSeries, a relationship that is created between a device and one or more channel paths during the processing of certain conditions. See also *implicit allegiance*, *contingent allegiance*, and *reserved allegiance*.

allocated storage

The space that is allocated to volumes but not yet assigned. Contrast with *assigned storage*.

American National Standards Institute (ANSI)

An organization of producers, consumers, and general interest groups that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States. (A)

Anonymous

In the DS Storage Manager, the label on an icon that represents all connections that are using fibre-channel adapters

between the storage unit and hosts but are not completely defined to the storage unit. See also *anonymous host*, *pseudo host*, and *access-any mode*.

anonymous host

Synonym for *pseudo host*. Contrast with *Anonymous* and *pseudo host*.

ANSI See *American National Standards Institute*.

APAR See *authorized program analysis report*. (GC)

API See *application programming interface*.

application programming interface

An interface that allows an application program that is written in a high-level language to use specific data or functions of the operating system or another program.

arbitrated loop

A fibre-channel topology that enables the interconnection of a set of nodes. See also *point-to-point connection* and *switched fabric*.

array An ordered collection, or group, of physical devices (disk drive modules) that is used to define logical volumes or devices. In the storage unit, an array is a group of disks that the user designates to be managed by the RAID technique. See also *redundant array of independent disks*.

ASCII (American National Standard Code for Information Interchange) The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity check), that is used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A) Some organizations, including IBM, have used the parity bit to expand the basic code set.

assigned storage

The space that is allocated to a volume and that is assigned to a port.

authorized program analysis report (APAR)

A request for correction of a defect in a current release of an IBM-supplied program. (GC)

availability

The degree to which a system or resource is capable of performing its normal function. See *data availability*.

B

bay The physical space that is used for installing SCSI, ESCON, and fibre-channel host adapter cards. The DS8000 has four bays, two in each cluster. See also *service boundary*.

bit 1) Either of the digits 0 or 1 when used in the binary numeration system. (T)
2) The storage medium required to store a single binary digit. See also *byte*.

block A string of data elements recorded or transmitted as a unit. The elements may be characters, words, or physical records. (T)

A group of consecutive bytes used as the basic storage unit in fixed-block architecture (FBA). All blocks on the storage device are the same size (fixed size). See also *fixed-block architecture* and *data record*.

byte 1) A group of eight adjacent binary digits that represent one EBCDIC character.
2) The storage medium required to store eight bits. See also *bit*.

C

cache A special-purpose buffer storage, smaller and faster than main storage, used to hold a copy of instructions and data obtained from main storage and likely to be needed next by the processor. (T)

cache fast write

A form of the fast-write operation in which the storage server writes the data directly to cache, where it is available for later destaging.

cache hit

An event that occurs when a read operation is sent to the cluster, and the requested data is found in cache. Contrast with *cache miss*.

cache memory

Memory, typically volatile memory, that a storage server uses to improve access times to instructions or data. The cache

memory is typically smaller and faster than the primary memory or storage medium. In addition to residing in cache memory, the same data also resides on the storage devices in the storage unit.

cache miss

An event that occurs when a read operation is sent to the cluster, but the data is not found in cache. Contrast with *cache hit*.

call home

A communication link established between the DS8000 and a service provider. The DS8000 can use this link to place a call to IBM or to another service provider when it requires service. With access to the machine, service personnel can perform service tasks, such as viewing error logs and problem logs or initiating trace and dump retrievals. See also *heartbeat* and *remote technical assistance information network*.

cascading

1) Connecting network controllers to each other in a succession of levels to concentrate many more lines than a single level permits.

2) In high-availability cluster multiprocessing (HACMP), cascading pertains to a cluster configuration in which the cluster node with the highest priority for a particular resource acquires the resource if the primary node fails. The cluster node relinquishes the resource to the primary node upon reintegration of the primary node into the cluster.

catcher

A server that service personnel use to collect and retain status data that an DS8000 sends to it.

CCR See *channel command retry*.

CCW See *channel command word*.

CD See *compact disc*.

CEC See *computer-electronic complex*.

channel

For zSeries, the part of a channel subsystem that manages a single I/O interface between a channel subsystem and a set of control units.

channel command retry (CCR)

For zSeries, the protocol used between a channel and a control unit that enables the control unit to request that the channel reissue the current command.

channel command word (CCW)

For zSeries, a data structure that specifies an I/O operation to the channel subsystem.

channel path

For zSeries, the interconnection between a channel and its associated control units.

channel subsystem

For zSeries, the part of a host computer that manages I/O communication between the program and any attached control units.

channel-subsystem image

For zSeries, the logical functions that a system requires to perform the function of a channel subsystem. With ESCON multiple image facility (EMIF), one channel subsystem image exists in the channel subsystem for each logical partition (LPAR). Each image appears to be an independent channel subsystem program, but all images share a common set of hardware facilities.

CKD See *count key data*.

CLI See *command-line interface*. See also *Copy Services command-line interface*.

cluster

1) A partition capable of performing all DS8000 functions. With two clusters in the DS8000, any operational cluster can take over the processing of a failing cluster.

2) In the AIX operating system, a group of nodes within a complex.

cluster processor complex (CPC)

The unit within a cluster that provides the management function for the DS8000. It consists of cluster processors, cluster memory, and related logic.

command-line interface (CLI)

An interface provided by an operating system that defines a set of commands and enables a user (or a script-like language) to issue these commands by typing text in response to the command prompt (for example, DOS commands or

UNIX shell commands). See also *Copy Services command-line interface*.

compact disc

An optically read disc, typically storing approximately 660 MB. CD-ROM (compact disc read-only memory) refers to the read-only format used to distribute DS8000 code and documentation.

compression

1) The process of eliminating gaps, empty fields, redundancies, and unnecessary data to shorten the length of records or blocks.

2) Any encoding that reduces the number of bits used to represent a given message or record. (GC)

computer-electronic complex (CEC)

The set of hardware facilities associated with a host computer.

concurrent copy

A facility on a storage server that enables a program to make a backup of a data set while the logical volume remains available for subsequent processing. The data in the backup copy is frozen at the point in time that the server responds to the request.

concurrent installation of licensed internal code

Process of installing licensed internal code on a DS8000 while applications continue to run.

concurrent maintenance

Service that is performed on a unit while it is operational.

concurrent media maintenance

Service performed on a disk drive module (DDM) without losing access to the data.

configure

In storage, to define the logical and physical configuration of the input/output (I/O) subsystem through the user interface that the storage unit provides for this function.

consistency group

A group of volumes participating in FlashCopy relationships in a logical subsystem, across logical subsystems, or across multiple storage units that must be kept in a consistent state to ensure data integrity.

consistency group interval time

The value in seconds that indicates the length of time between the formation of consistency groups.

consistent copy

A copy of a data entity (a logical volume, for example) that contains the contents of the entire data entity at a single instant in time.

console

A user interface to a server, for example, the interface provided on a personal computer. See also *IBM TotalStorage Management Console*.

contingent allegiance

For zSeries, a relationship that is created in a control unit between a device and a channel when the channel accepts unit-check status. The allegiance causes the control unit to guarantee access; the control unit does not present the busy status to the device. The allegiance enables the channel to retrieve sense data that is associated with the unit-check status on the channel path associated with the allegiance.

control path

The route that is established from the master storage unit to the subordinate storage unit when more than one storage unit participates in a Global Mirror session. If there is only one storage unit (the master) in the Global Mirror session, no control path is required.

control unit (CU)

- 1) A device that coordinates and controls the operation of one or more input/output devices, and synchronizes the operation of such devices with the operation of the system as a whole.
- 2) For zSeries, a storage server with ESCON or OEMI interfaces. The control unit adapts a native device interface to an I/O interface that a zSeries host system supports.
- 3) The portion of the storage unit that supports the attachment of emulated count key data devices over ESCON, FICON, or OEMI interfaces. See also *cluster*.

control-unit image

For zSeries, a logical subsystem that is

accessed through an ESCON I/O interface. One or more control-unit images exist in each control unit. Each image appears as an independent control unit, but all control-unit images share a common set of hardware facilities. The DS8000 can emulate 3990-3, TPF, 3990-6, or 2105 control units.

control-unit-initiated reconfiguration (CUIR)

A software mechanism that the DS8000 uses to request that an operating system of a zSeries host verify that one or more subsystem resources can be taken offline for service. The DS8000 can use this process to automatically vary channel paths offline and online to facilitate bay service or concurrent code installation. Depending on the operating system, support for this process might be model dependent, might depend on the IBM TotalStorage Enterprise Storage Server Subsystem Device Driver, or might not exist.

Coordinated Universal Time (UTC)

The international standard of time that is kept by atomic clocks around the world.

Copy Services

A collection of optional software features, with a Web-browser interface, used for configuring, managing, and monitoring data-copy functions.

Copy Services CLI

See *Copy Services command-line interface*.

Copy Services client

Software that runs on each DS8000 cluster in the Copy Services server group and that performs the following functions:

- Communicates configuration, status and connectivity information to the Copy Services server
- Performs data-copy functions on behalf of the Copy Services server

Copy Services command-line interface (Copy Services CLI)

The command-line interface software that is provided with DS8000 Copy Services and used for invoking Copy Services functions from host systems attached to the DS8000. See also *command-line interface*.

Copy Services server

A cluster that the Copy Services administrator designates to perform the DS8000 Copy Services functions.

Copy Services server group

A collection of user-designated DS8000 clusters participating in Copy Services functions that a designated, active, Copy Services server manages. A Copy Services server group is also called a Copy Services domain.

count field

The first field of a count key data (CKD) record. This eight-byte field contains a four-byte track address (CCHH). It defines the cylinder and head that are associated with the track, and a one-byte record number (R) that identifies the record on the track. It defines a one-byte key length that specifies the length of the record's key field (0 means no key field). It defines a two-byte data length that specifies the length of the record's data field (0 means no data field). Only the end-of-file record has a data length of zero.

count key data (CKD)

For zSeries, a data-record format employing self-defining record formats in which each record is represented by up to three fields: a *count* field that identifies the record and specifies its format, an optional *key* field that identifies the data area contents, and an optional *data* field that typically contains the user data. For CKD records on the storage unit, the logical volume size is defined in terms of the device emulation mode (3390 or 3380 track format). The count field is always 8 bytes long and contains the lengths of the key and data fields, the key field has a length of 0 to 255 bytes, and the data field has a length of 0 to 65 535 or the maximum that will fit on the track. See also *data record*.

CPC See *cluster processor complex*.

CRC See *cyclic redundancy check*.

CU See *control unit*.

CUIR See *control-unit initiated reconfiguration*.

custom volume

A volume in count-key-data (CKD) format that is not a standard volume, which means that it does not necessarily present

the same number of cylinders and capacity to its assigned logical control unit as provided by one of the following standard zSeries volume types: 3390-2, 3390-3, 3390-9, 3390-2 (3380-track mode), or 3390-3 (3380-track mode). See also *count-key-data*, *interleave*, *standard volume*, and *volume*.

CUT See *Coordinated Universal Time*.

cyclic redundancy check (CRC)

A redundancy check in which the check key is generated by a cyclic algorithm. (T)

cylinder

A unit of storage on a CKD device with a fixed number of tracks.

D

DA See *device adapter*. See also *IBM Serial Storage adapter*.

daisy chain

See *serial connection*.

DASD See *direct access storage device*.

DASD fast write (DFW)

A function of a storage server in which active write data is stored in nonvolatile cache, thus avoiding exposure to data loss.

data availability

The degree to which data is available when needed, typically measured as a percentage of time that the system would be capable of responding to any data request (for example, 99.999% available).

data compression

A technique or algorithm used to encode data such that the encoded result can be stored in less space than the original data. The original data can be recovered from the encoded result through a reverse technique or reverse algorithm. See also *compression*.

Data Facility Storage Management Subsystem (DFSMS)

An operating environment that helps automate and centralize the management of storage. To manage storage, DFSMS provides the storage administrator with control over data class, storage class, management class, storage group, and automatic class selection routine definitions.

data field

The optional third field of a count key data (CKD) record. The count field specifies the length of the data field. The data field contains data that the program writes.

data record

The basic unit of zSeries storage on a DS8000, also known as a count-key-data (CKD) record. Data records are stored on a track. The records are sequentially numbered starting with 0. The first record, R0, is typically called the track descriptor record and contains data that the operating system normally uses to manage the track. See also *count-key-data* and *fixed-block architecture*.

data set FlashCopy

For zSeries hosts, a feature of FlashCopy that indicates how many partial volume FlashCopy relationships are active on a volume.

data sharing

The ability of multiple host systems to concurrently utilize data that they store on one or more storage devices. The storage unit enables configured storage to be accessible to any, or all, attached host systems. To use this capability, the host program must be designed to support data that it is sharing.

DDM See *disk drive module*.

DDM group

See *disk pack*.

dedicated storage

Storage within a storage unit that is configured such that a single host system has exclusive access to the storage.

demote

To remove a logical data unit from cache memory. A storage server demotes a data unit to make room for other logical data units in the cache or because the logical data unit is not valid. The storage unit must destage logical data units with active write units before they can be demoted. See also *destage*.

destage

To move data from an online or higher priority to an offline or lower priority

device. The storage unit stages incoming data into cache and then destages it to disk.

device

For zSeries, a disk drive.

device adapter (DA)

A physical component of the DS8000 that provides communication between the clusters and the storage devices. The DS8000 has eight device adapters that it deploys in pairs, one from each cluster. Device adapter pairing enables the DS8000 to access any disk drive from either of two paths, providing fault tolerance and enhanced availability.

device address

For zSeries, the field of an ESCON device-level frame that selects a specific device on a control-unit image.

device ID

The unique two-digit hexadecimal number that identifies the logical device.

device interface card

A physical subunit of a storage cluster that provides the communication with the attached device drive modules.

device number

For zSeries, a four-hexadecimal-character identifier, for example 13A0, that the systems administrator associates with a device to facilitate communication between the program and the host operator. The device number is associated with a subchannel.

device sparing

A subsystem function that automatically copies data from a failing device drive module to a spare device drive module. The subsystem maintains data access during the process.

DFS See *distributed file service*.

DFSMS

See *Data Facility Storage Management Subsystem*.

direct access storage device (DASD)

1) A mass storage medium on which a computer stores data.

2) A disk device.

disk cage

A container for disk drives. Each disk cage supports eight disk packs (64 disks).

disk drive

Standard term for a disk-based nonvolatile storage medium. The DS8000 uses hard disk drives as the primary nonvolatile storage media to store host data.

disk drive module (DDM)

A field replaceable unit that consists of a single disk drive and its associated packaging.

disk drive module group

See *disk pack*.

disk pack

A group of disk drive modules (DDMs) installed as a unit in a DDM bay.

disk group

A collection of 4 disk drives that are connected to the same pair of IBM Serial Storage adapters and can be used to create a RAID array. A disk group can be formatted as count key data or fixed block, and as RAID or non-RAID, or it can be left unformatted. A disk group is a logical assemblage of disk drives. Contrast with *disk pack*.

distributed file service (DFS)

A service that provides data access over IP networks.

DNS See *domain name system*.

domain

- 1) That part of a computer network in which the data processing resources are under common control.
- 2) In TCP/IP, the naming system used in hierarchical networks.
- 3) A Copy Services server group, in other words, the set of clusters the user designates to be managed by a particular Copy Services server.

domain name system (DNS)

In TCP/IP, the server program that supplies name-to-address translation by mapping domain names to internet addresses. The address of a DNS server is the internet address of the server that hosts the DNS software for the network.

dotted decimal notation

A convention used to identify IP

addresses. The notation consists of four 8-bit numbers written in base 10. For example, 9.113.76.250 is an IP address that contains the octets 9, 113, 76, and 250.

drawer

A unit that contains multiple device drive modules and provides power, cooling, and related interconnection logic to make the device drive modules accessible to attached host systems.

drive 1) A peripheral device, especially one that has addressed storage media. See also *disk drive module*.

2) The mechanism used to seek, read, and write information on a storage medium.

DS8000

See *IBM TotalStorage DS8000*.

DS8000 Batch Configuration tool

A program that automatically configures a DS8000. The configuration is based on data that IBM service personnel enter into the program.

DS Storage Manager

See *IBM TotalStorage DS Storage Manager*.

duplex

- 1) Regarding Copy Services, the state of a volume pair after Remote Mirror and Copy has completed the copy operation and the volume pair is synchronized.
- 2) In general, pertaining to a communication mode in which data can be sent and received at the same time.

dynamic sparing

The ability of a storage server to move data from a failing disk drive module (DDM) to a spare DDM while maintaining storage functions.

E

E10 The predecessor of the F10 model of the Enterprise Storage Server. See also *F10*.

E20 The predecessor of the F20 model of the Enterprise Storage Server. See also *F20*.

EBCDIC

See *extended binary-coded decimal interchange code*.

EC See *engineering change*.

ECKD See *extended count key data*.

eight pack
See *disk pack*.

electrostatic discharge (ESD)
An undesirable discharge of static electricity that can damage equipment and degrade electrical circuitry.

emergency power off (EPO)
A means of turning off power during an emergency, usually a switch.

EMIF See *ESCON multiple image facility*.

enclosure
A unit that houses the components of a storage subsystem, such as a control unit, disk drives, and power source.

end of file
A coded character recorded on a data medium to indicate the end of the medium. On a count-key-data direct access storage device, the subsystem indicates the end of a file by including a record with a data length of zero.

engineering change (EC)
An update to a machine, part, or program.

Enterprise Systems Architecture/390 (ESA/390)
An IBM architecture for mainframe computers and peripherals. Processor systems that follow the ESA/390 architecture include the ES/9000® family. See also *z/Architecture*.

Enterprise Systems Connection (ESCON)
1) A zSeries computer peripheral interface. The I/O interface uses zSeries logical protocols over a serial interface that configures attached units to a communication fabric.
2) A set of IBM products and services that provide a dynamically connected environment within an enterprise.

EPO See *emergency power off*.

ERDS See *error-recording data set*.

ERP See *error recovery procedure*.

error-recording data set (ERDS)
On zSeries hosts, a data set that records data-storage and data-retrieval errors. A

service information message (SIM) provides the error information for the ERDS.

error recovery procedure (ERP)
Procedures designed to help isolate and, where possible, to recover from errors in equipment. The procedures are often used in conjunction with programs that record information on machine malfunctions.

ESA/390
See *Enterprise Systems Architecture/390*.

ESCD See *ESCON director*.

ESCON
See *Enterprise System Connection*.

ESCON channel
A zSeries channel that supports ESCON protocols.

ESCON director (ESCD)
An I/O interface switch that allows the interconnection of multiple ESCON interfaces in a distributed-star topology.

ESCON host systems
zSeries hosts that attach to the DS8000 with an ESCON adapter. Such host systems run on operating systems that include MVS, VSE, TPF, or versions of VM.

ESCON multiple image facility (EMIF)
For zSeries, a function that enables LPARs to share an ESCON channel path by providing each LPAR with its own channel-subsystem image.

EsconNet
In the DS Storage Manager, the label on a pseudo host icon that represents a host connection that uses the ESCON protocol and that is not completely defined on the DS8000. See also *pseudo host* and *access-any mode*.

ESD See *electrostatic discharge*.

eserver
See *IBM e(logo)server*.

ESSNet
See *IBM TotalStorage Enterprise Storage Server Network*.

extended binary-coded decimal interchange code (EBCDIC)
An IBM-developed coding scheme used to represent various alphabetic, numeric, and

special symbols with a coded character set of 256 eight-bit codes.

extended count key data (ECKD)

An extension of the count key data (CKD) architecture.

extent A continuous space on a disk that is occupied by or reserved for a particular data set, data space, or file. The unit of increment is a track. See also *multiple allegiance* and *parallel access volumes*.

extent pool

A groups of extents. See also *extent*.

F

fabric In fibre-channel technology, a routing structure, such as a switch, receives addressed information and routes to the appropriate destination. A fabric can consist of more than one switch. When multiple fibre-channel switches are interconnected, they are said to be *cascaded*.

failback

Pertaining to a cluster recovery from failover following repair. See also *failover*.

failover

Pertaining to the process of transferring all control to a single cluster when the other cluster in the storage unit fails. See also *cluster* and *failback*.

fast write

A write operation at cache speed that does not require immediate transfer of data to a disk drive. The subsystem writes the data directly to cache, to nonvolatile storage, or to both. The data is then available for destaging. A fast-write operation reduces the time an application must wait for the I/O operation to complete.

FBA See *fixed-block architecture*.

FC See *feature code*. **Note:** *FC* is a common abbreviation for fibre channel in the industry, but the DS8000 customer documentation library reserves *FC* for feature code.

FC-AL See *Fibre Channel-Arbitrated Loop*.

FCP See *fibre-channel protocol*.

FCS See *fibre-channel standard*.

feature code (FC)

A code that identifies a particular orderable option and that is used by service personnel to process hardware and software orders. Individual optional features are each identified by a unique feature code.

fibre channel

A data-transmission architecture based on the ANSI Fibre Channel standard, which supports full-duplex communication. The DS8000 supports data transmission over fiber-optic cable through its fibre-channel adapters. See also *fibre-channel protocol* and *fibre-channel standard*.

Fibre Channel-Arbitrated Loop (FC-AL)

An implementation of the Fibre Channel standard that uses a ring topology for the communication fabric. Refer to American National Standards Institute (ANSI) X3T11/93-275. In this topology, two or more fibre-channel end points are interconnected through a looped interface.

fibre-channel connection (FICON)

A fibre-channel communications protocol designed for IBM mainframe computers and peripherals.

fibre-channel protocol (FCP)

A protocol used in fibre-channel communications with five layers that define how fibre-channel ports interact through their physical links to communicate with other ports.

Fibre-Channel standard (FCS)

An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. The protocol has two layers. The IP layer defines basic interconnection protocols. The upper layer supports one or more logical protocols (for example, FCP for SCSI command protocols and SBICON for zSeries command protocols). Refer to American National Standards Institute (ANSI) X3.230-199x. See also *fibre-channel protocol*.

fibre-channel topology

An interconnection topology supported on fibre-channel adapters. See also *point-to-point connection*, *switched fabric*, and *arbitrated loop*.

FICON

See *fibre-channel connection*.

FiconNet

In the DS Storage Manager, the label on a pseudo host icon that represents a host connection that uses the FICON protocol and that is not completely defined on the DS8000. See also *pseudo host* and *access-any mode*.

field replaceable unit (FRU)

An assembly that is replaced in its entirety when any one of its components fails. In some cases, a field replaceable unit might contain other field replaceable units. (GC)

FIFO See *first-in-first-out*.

File Transfer Protocol (FTP)

In TCP/IP, an application protocol used to transfer files to and from host computers. See also *Transmission Control Protocol/Internet Protocol*.

firewall

A protection against unauthorized connection to a computer or a data storage system. The protection is usually in the form of software on a gateway server that grants access to users who meet authorization criteria.

first-in-first-out (FIFO)

A queuing technique in which the next item to be retrieved is the item that has been in the queue for the longest time. (A)

fixed-block architecture (FBA)

An architecture for logical devices that specifies the format of and access mechanisms for the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size). The subsystem can access them independently.

fixed-block device

An architecture for logical devices that specifies the format of the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size); the subsystem can access them independently. This is the required format of the logical data units for host systems that attach with a SCSI or fibre-channel interface. See also *fibre channel* and *small computer systems interface*.

FlashCopy

An optional feature of the DS8000 that can make an instant copy of data, that is, a point-in-time copy of a volume.

FlashCopy relationship

A mapping of a FlashCopy source volume and a FlashCopy target volume that allows a point-in-time copy of the source volume to be copied to the target volume. FlashCopy relationships exist from the time that you initiate a FlashCopy operation until the storage unit copies all data from the source volume to the target volume or until you delete the FlashCopy relationship, if it is persistent.

FRU See *field replaceable unit*.

FTP See *File Transfer Protocol*.

full duplex

See *duplex*.

fuzzy copy

A function of the Global Copy feature wherein modifications to the primary logical volume are performed on the secondary logical volume at a later time. The original order of update is not strictly maintained. See also *Global Copy*.

G

GB See *gigabyte*.

GDPS See *Geographically Dispersed Parallel Sysplex*.

Geographically Dispersed Parallel Sysplex (GDPS)

A zSeries multisite application-availability solution.

gigabyte (GB)

A gigabyte of storage is 10^9 bytes. A gigabyte of memory is 2^{30} bytes.

Global Copy

An optional capability of the DS8000 remote mirror and copy feature that maintains a fuzzy copy of a logical volume on the same DS8000 or on another DS8000. In other words, all modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume at a later point in time. The original order of update is not strictly maintained. See also *Remote Mirror and Copy* and *Metro Mirror*.

Global Mirror

An optional capability of the DS8000 remote mirror and copy feature that provides a 2-site extended distance remote copy. Data that is written by the host to the storage unit at the local site is automatically maintained at the remote site. See also *Metro Mirror* and *Remote Mirror and Copy*.

group In DS8000 documentation, a nickname for two different kinds of groups, depending on the context. See *disk pack* or *Copy Services server group*.

H

HA See *host adapter*.

HACMP

See *High-Availability Cluster Multi-Processing*.

hard disk drive (HDD)

1) A storage medium within a storage server used to maintain information that the storage server requires.

2) A mass storage medium for computers that is typically available as a fixed disk (such as the disks used in system units of personal computers or in drives that are external to a personal computer) or a removable cartridge.

hardware service manager (HSM)

An option on an AS/400 or iSeries host that enables the user to display and work with system hardware resources and to debug input-output processors (IOP), input-output adapters (IOA), and devices.

HCD See *Hardware Configuration Data*.

HDA See *head and disk assembly*.

HDD See *hard disk drive*.

hdisk An AIX term for storage space.

head and disk assembly (HDA)

The portion of an HDD associated with the medium and the read/write head.

heartbeat

A status report sent at regular intervals from the DS8000. The service provider uses this report to monitor the health of the call home process. See also *call home*, *heartbeat call home record*, and *remote technical assistance information network*.

heartbeat call home record

Machine operating and service information sent to a service machine. These records might include such information as feature code information and product logical configuration information.

hierarchical storage management

1) A function in storage management software, such as Tivoli Storage Management or Data Facility Storage Management Subsystem/MVS (DFSMS/MVS), that automatically manages free space based on the policy that the storage administrator sets.

2) In AS/400 storage management, an automatic method to manage and distribute data between the different storage layers, such as disk units and tape library devices.

High-Availability Cluster Multi-Processing (HACMP)

Software that provides host clustering, so that a failure of one host is recovered by moving jobs to other hosts within the cluster.

high-speed link (HSL)

A hardware connectivity architecture that links system processors to system input/output buses and other system units.

home address

A nine-byte field at the beginning of a track that contains information that identifies the physical track and its association with a cylinder.

hop Interswitch connection. A hop count is the number of connections that a particular block of data traverses between source and destination. For example, data traveling from one hub over a wire to another hub traverses one hop.

host See *host system*.

host adapter

A physical subunit of a storage server that provides the ability to attach to one or more host I/O interfaces.

host name

The Internet address of a machine in the network. The host name can be entered in the host definition as the fully qualified domain name of the attached host system, such as *mycomputer.city.company.com*, or

as the subname of the fully qualified domain name, for example, mycomputer. See also *host system*.

host processor

A processor that controls all or part of a user application network. In a network, the processing unit in which the data communication access method resides. See also *host system*.

host system

A computer, either of the mainframe (for example, zSeries) or of the open-systems type, that is connected to the DS8000. zSeries hosts are connected to the DS8000 through ESCON interfaces. Open-systems hosts are connected to the DS8000 by SCSI or fibre-channel interfaces.

hot plug

Pertaining to the ability to add or remove a hardware facility or resource to a unit while power is on.

HSL See *high-speed link*.

HSM See *hierarchical storage management or Hardware Service Manager*.

I

i5/OS The IBM operating system that runs the IBM i5/OS and e(logo)Server i5 server families of servers.

IBM e(logo)server

The IBM brand name for a series of server products that are optimized for e-commerce. The products include the iSeries, pSeries, xSeries, and zSeries.

IBM product engineering (PE)

The third-level of IBM service support. Product engineering is composed of IBM engineers who have experience in supporting a product or who are knowledgeable about the product.

IBM Serial Storage adapter

A physical adapter based on the IBM Serial Storage architecture. IBM Serial Storage adapters connect disk drive modules to DS8000 clusters.

IBM TotalStorage

The brand name used to identify storage products from IBM, including the IBM TotalStorage DS8000. See also *IBM*

TotalStorage DS8000 and IBM TotalStorage DS Storage Manager.

IBM TotalStorage DS8000

A member of the IBM TotalStorage Resiliency Family of storage servers and attached storage devices (disk drive modules). The DS8000 delivers high-performance, fault-tolerant storage and management of enterprise data, affording access through multiple concurrent operating systems and communication protocols. High performance is provided by multiple symmetrical multiprocessors, integrated caching, RAID support for the disk drive modules, and disk access through a high-speed serial storage architecture interface.

IBM TotalStorage DS CLI

The command-line interface (CLI) that is specific to the DS8000.

IBM TotalStorage DS Storage Manager (DS Storage Manager)

Software with a Web-browser interface for configuring the DS8000.

IBM TotalStorage Enterprise Storage Server Network (ESSNet)

A private network providing Web browser access to the Enterprise Storage Server. IBM installs the ESSNet software on an IBM workstation called the IBM TotalStorage ESS Master Console, supplied with the first ESS delivery.

IBM TotalStorage Management Console (MC)

An IBM workstation that acts as the focal point for configuration, Copy Services management, and maintenance for the DS8000. It includes a Web browser that provides links to the user interface, including the DS Storage Manager and the DS8000 Copy Services.

IBM TotalStorage Multipath Subsystem Device Driver (SDD)

Software that is designed to support the multipath configuration environments of the DS8000. The SDD resides in a host system with the native disk device driver.

IBM TotalStorage Resiliency Family

A set of hardware and software features and products, as well as integrated software and services that are available on the IBM TotalStorage DS8000 and the

IBM TotalStorage Enterprise Storage Server, Models 750 and 800.

image See *storage image*.

IML See *initial microcode load*.

implicit allegiance

In Enterprise Systems Architecture/390, a relationship that a control unit creates between a device and a channel path when the device accepts a read or write operation. The control unit guarantees access to the channel program over the set of channel paths that it associates with the allegiance.

initial microcode load (IML)

The action of loading microcode for a computer into that computer's storage.

initial program load (IPL)

The action of loading software into a computer, typically an operating system that controls the computer.

initiator

A SCSI device that communicates with and controls one or more targets. Contrast with *target*.

i-node The internal structure in an AIX operating system that describes the individual files in the operating system. It contains the code, type, location, and owner of a file.

input/output (I/O)

Pertaining to (a) input, output, or both or (b) a device, process, or channel involved in data input, data output, or both.

input/output configuration data set

A configuration definition built by the I/O configuration program (IOCP) and stored on disk files associated with the processor controller.

interleave

To automatically create two striped partitions across the drives in a RAID-5 array, both of which use the count-key-data (CKD) record format.

Internet Protocol (IP)

In the Internet suite of protocols, a protocol without connections that routes data through a network or interconnecting networks and acts as an intermediary between the higher protocol layers and the physical network. The upper layer supports one or more logical protocols (for

example, a SCSI-command protocol and a zSeries command protocol). Refer to ANSI X3.230-199x. The IP acronym is the IP in TCP/IP. See also *Transmission Control Protocol/Internet Protocol*.

invalidate

To remove a logical data unit from cache memory because it cannot support continued access to the logical data unit on the device. This removal might be the result of a failure within the storage server or a storage device that is associated with the device.

I/O See *input/output*.

I/O adapter (IOA)

An input-output adapter on the PCI bus.

IOCDS

See *input/output configuration data set*.

IOCP See *I/O Configuration Program*.

I/O Configuration Program (IOCP)

A program that defines to a system all the available I/O devices and channel paths.

I/O device

An addressable read and write unit, such as a disk drive device, magnetic tape device, or printer.

I/O interface

An interface that enables a host to perform read and write operations with its associated peripheral devices.

I/O Priority Queueing

A facility in the Workload Manager of zSeries that enables the system administrator to set priorities for queueing I/Os from different system images. See also *multiple allegiance* and *parallel access volume*.

I/O processor (IOP)

Controls input-output adapters and other devices.

I/O sequential response time

The time an I/O request is queued in processor memory waiting for previous I/Os to the same volume to complete.

IOSQ See *I/O sequential response time*.

IP See *Internet Protocol*.

IPL See *initial program load*.

iSeries

An IBM e(logo)server product that emphasizes integration. It is the successor to the AS/400 family of servers.

J

Java Virtual Machine (JVM)

A software implementation of a central processing unit (CPU) that runs compiled Java code (applets and applications). (GC)

JVM See *Java Virtual Machine*.

K

KB See *kilobyte*.

key field

The second (optional) field of a count key data record. The key length is specified in the count field. The key length determines the field length. The program writes the data in the key field and uses the key field to identify or locate a given record. The subsystem does not use the key field.

kilobyte (KB)

1) For processor storage, real, and virtual storage, and channel volume, 2^{10} or 1024 bytes.

2) For disk storage capacity and communications volume, 1000 bytes.

Korn shell

Interactive command interpreter and a command programming language.

KPOH See *thousands of power-on hours*.

L

LAN See *local area network*.

last-in first-out (LIFO)

A queuing technique in which the next item to be retrieved is the item most recently placed in the queue. (A)

LBA See *logical block address*.

LCU See *logical control unit*.

least recently used (LRU)

1) The algorithm used to identify and make available the cache space that contains the least-recently used data.

2) A policy for a caching algorithm that chooses to remove from cache the item that has the longest elapsed time since its last access.

LED See *light-emitting diode*.

LMC See *licensed machine code*.

licensed machine code (LMC)

Microcode that IBM does not sell as part of a machine, but licenses to the customer. LMC is implemented in a part of storage that is not addressable by user programs. Some IBM products use it to implement functions as an alternate to hard-wired circuitry.

LIFO See *last-in first-out*.

light-emitting diode (LED)

A semiconductor chip that gives off visible or infrared light when activated.

link address

On an ESCON interface, the portion of a source or destination address in a frame that ESCON uses to route a frame through an ESCON director. ESCON associates the link address with a specific switch port that is on the ESCON director. Equivalently, it associates the link address with the channel subsystem or control unit link-level functions that are attached to the switch port.

link-level facility

The ESCON hardware and logical functions of a control unit or channel subsystem that allow communication over an ESCON write interface and an ESCON read interface.

local area network (LAN)

A computer network located on a user's premises within a limited geographic area.

local e-mail

An e-mail configuration option for storage servers that are connected to a host-system network that does not have a domain name system (DNS) server.

logical address

On an ESCON interface, the portion of a source or destination address in a frame used to select a specific channel-subsystem or control-unit image.

logical block address (LBA)

The address assigned by the DS8000 to a sector of a disk.

logical control unit (LCU)

See *control-unit image*.

logical data unit

A unit of storage that is accessible on a given device.

logical device

The facilities of a storage server (such as the DS8000) associated with the processing of I/O operations directed to a single host-accessible emulated I/O device. The associated storage is referred to as a logical volume. The logical device is mapped to one or more host-addressable units, such as a device on a zSeries I/O interface or a logical unit on a SCSI I/O interface, such that the host initiating I/O operations to the I/O-addressable unit interacts with the storage on the associated logical device.

logical partition (LPAR)

For zSeries, a set of functions that create the programming environment in which more than one logical partition (LPAR) is established on a processor. An LPAR is conceptually similar to a virtual machine environment except that the LPAR is a function of the processor. Also, the LPAR does not depend on an operating system to create the virtual machine environment.

logical path

1) The relationship between a channel image and a control-unit image that designates the physical path to be used for device-level communications between these images. The logical path is established as part of the channel and control-unit initialization procedures by the exchange of link-level frames.

2) With the Remote Mirror and Copy feature, the relationship between a source logical subsystem (LSS) and a target LSS that is created over a physical path through the interconnection fabric that is used for Remote Mirror and Copy functions. An LSS is a primary control unit, which performs the functions of a channel image.

logical subsystem (LSS)

A topological construct that consists of a

group of up to 256 logical devices. A DS8000 can have up to 16 CKD-formatted logical subsystems (4096 CKD logical devices) and also up to 16 fixed-block logical subsystems (4096 fixed-block logical devices). The logical subsystem facilitates configuration of the DS8000 and might have other implications relative to the operation of certain functions. There is a one-to-one mapping between a CKD logical subsystem and a zSeries control-unit image.

For zSeries hosts, a logical subsystem represents a logical control unit (LCU). Each control-unit image is associated with only one logical subsystem. See also *control-unit image*.

logical unit

In open systems, a logical disk drive.

logical unit number (LUN)

In the SCSI protocol, a unique number that is used on a SCSI bus to enable it to differentiate between separate devices, each of which is a logical unit.

logical volume

The storage medium that is associated with a logical disk drive. A logical volume typically resides on one or more storage devices. The DS8000 administrator defines this unit of storage. The logical volume, when residing on a RAID-formatted array, is spread over the drives in the array.

logical volume manager (LVM)

A set of system commands, library routines, and other tools that allow the user to establish and control logical volume storage. The LVM maps data between the logical view of storage space and the physical disk drive module.

longitudinal redundancy check (LRC)

1) A method of error checking during data transfer that involves checking parity on a row of binary digits that are members of a set that forms a matrix. Longitudinal redundancy check is also called a longitudinal parity check.

2) A mechanism that the DS8000 uses for locating errors. The LRC checks the data as it progresses from the host, through the DS8000 controller, into the device adapter, and to the array.

longwave laser adapter

A connector that is used between a host and the DS8000 to support longwave fibre-channel communication.

loop The physical connection between a pair of device adapters in the DS8000. See also *device adapter*.

LPAR See *logical partition*.

LRC See *longitudinal redundancy check*.

LRU See *least recently used*.

LSS See *logical subsystem*.

LUN See *logical unit number*.

LVM See *logical volume manager*.

M**machine level control (MLC)**

A database that contains the EC level and configuration of products in the field.

machine reported product data (MRPD)

Product data gathered by a machine and sent to a destination such as an IBM support server or RETAIN. These records might include such information as feature code information and product logical configuration information.

mainframe

A computer, usually in a computer center, with extensive capabilities and resources to which other computers may be connected so that they can share facilities. (T)

maintenance analysis procedure (MAP)

A hardware maintenance document that gives an IBM service representative a step-by-step procedure for tracing a symptom to the cause of a failure.

management console

See *IBM TotalStorage Management Console*.

Management Information Base (MIB)

1) A collection of objects that can be accessed by means of a network management protocol. (GC)

2) The MIB record conforms to the Open Systems Interconnection (OSI) standard defined by the International Organization for Standardization (ISO) for the exchange of information. See also *simple network management protocol*.

MAP See *maintenance analysis procedure*.

master storage unit

The physical unit that controls the creation of consistency groups in a Global Mirror session. The master storage unit sends commands to subordinate storage units. A storage unit can be a master for only one Global Mirror session. Contrast with *subordinate storage unit*.

maximum consistency group drain time

The value in seconds that indicates the maximum time that writes from the local site are delayed to the remote site while the current consistency group is being formed at the remote site. When this time is exceeded, the current attempt to form a consistency group is ended and another attempt is started. If this time is exceeded five times, this maximum time is ignored on the next attempt to form a consistency group. The default value is the larger of four minutes or two times the consistency group interval time if this value is set to zero.

maximum coordination time

The value in milliseconds that indicates the maximum time that is allowed for host I/O to be delayed during the coordination of the primary volumes of an Global Mirror session. The default is 50 milliseconds if this value is set to zero.

MB See *megabyte*.

MC See *IBM TotalStorage Management Console*.

MCA See *Micro Channel architecture*.

MDM See *Multiple Device Manager*.

mean time between failures (MTBF)

1) A projection of the time that an individual unit remains functional. The time is based on averaging the performance, or projected performance, of a population of statistically independent units. The units operate under a set of conditions or assumptions.

2) For a stated period in the life of a functional unit, the mean value of the lengths of time between consecutive failures under stated conditions. (I) (A)

medium

For a storage unit, the disk surface on which data is stored.

megabyte (MB)

- 1) For processor storage, real and virtual storage, and channel volume, 2^{20} or 1 048 576 bytes.
- 2) For disk storage capacity and communications volume, 1 000 000 bytes.

Metro Mirror

A function of a storage server that maintains a consistent copy of a logical volume on the same storage server or on another storage server. All modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume. See also *Remote Mirror and Copy* and *Global Copy*.

MES See *miscellaneous equipment specification*.

MIB See *management information base*.

Micro Channel architecture (MCA)

The rules that define how subsystems and adapters use the Micro Channel bus in a computer. The architecture defines the services that each subsystem can or must provide.

Microsoft Internet Explorer (MSIE)

Web browser software manufactured by Microsoft.

migration

The replacement of a system or subsystem with a different type of system or subsystem, such as replacing a SCSI host adapter with a fibre-channel host adapter. In the context of data migration regarding the DS8000, the transfer of data from one storage unit to another, such as from a 3390 to the DS8000.

MIH See *missing-interrupt handler*.

mirrored pair

Two units that contain the same data. The system refers to them as one entity.

mirroring

In host systems, the process of writing the same data to two disk units within the same auxiliary storage pool at the same time.

miscellaneous equipment specification (MES)

IBM field-installed change to a machine.

missing-interrupt handler (MIH)

An MVS and MVS/XA facility that tracks

I/O interrupts. MIH informs the operator and creates a record whenever an expected interrupt fails to occur before a specified elapsed time is exceeded.

MLC See *machine level control*.

mobile solutions terminal (MoST)

The mobile terminal used by service personnel.

mode conditioning patch cable

A cable that converts a single-mode signal from a longwave adapter into a light signal that is appropriate for multimode fibre. Another mode conditioning patch cable is required at the terminating end of the multimode fibre to convert the signal back to a single-mode signal for a longwave adapter.

Model 100

A 2105 Model 100, often simply referred to as a Mod 100, is an expansion enclosure for the Enterprise Storage Server. See also *2105*.

MoST See *mobile solutions terminal*.

MRPD See *machine reported product data*.

MSA See *multiport serial adapter*.

MSIE See *Microsoft Internet Explorer*.

MTBF See *mean time between failures*.

Multipath Subsystem Device Driver

See *IBM TotalStorage DS8000 Multipath Subsystem Device Driver*.

multiple allegiance

A DS8000 hardware function that is independent of software support. This function enables multiple system images to concurrently access the same logical volume on the DS8000 as long as the system images are accessing different extents. See also *extent* and *parallel access volumes*.

Multiple Device Manager (MDM)

A component of the IBM TotalStorage Productivity Center that allows administrators to configure, manage, and monitor the performance of SAN storage devices from a single console.

multiple relationship FlashCopy

An option of the DS8000 that creates backup copies from one source to multiple

targets by simultaneously establishing multiple FlashCopy relationships.

multiple virtual storage (MVS)

Implies MVS/390, MVS/XA, MVS/ESA, and the MVS element of the zSeries operating system.

multiplex

The action of transmitting simultaneously.

multipoint serial adapter (MSA)

An adapter on the IBM TotalStorage Management Console that has multiple ports to which aDS8000 can be attached.

multiprocessor

A computer that includes two or more processors that have common access to a main storage. For the DS8000, the multiprocessors operate in parallel.

MVS See *multiple virtual storage*.

N

name server

A server that stores names of the participating DS8000 clusters.

Netfinity

IBM Intel-processor-based server; predecessor to the IBM xSeries server.

Netscape Navigator

Web browser software manufactured by Netscape.

network manager

A program or group of programs that is used to monitor, manage, and diagnose the problems of a network. (GC)

node The unit that is connected in a fibre-channel network. A DS8000 is a node in a fibre-channel network.

non-RAID

A disk drive set up independently of other disk drives and not set up as part of a disk pack to store data using the redundant array of disks (RAID) data-striping methodology.

nonremovable medium

A recording medium that cannot be added to or removed from a storage device.

nonvolatile storage (NVS)

Memory that stores active write data to avoid data loss in the event of a power loss.

NVS See *nonvolatile storage*.

O

octet In Internet Protocol addressing, one of the four parts of a 32-bit integer presented in dotted decimal notation. See also *dotted decimal notation*.

OEMI See *original equipment manufacturer's information*.

open system

A system whose characteristics comply with standards made available throughout the industry and that therefore can be connected to other systems complying with the same standards. Applied to the DS8000, such systems are those hosts that connect to the DS8000 through SCSI or FCP protocols. See also *small computer system interface* and *fibre-channel protocol*.

organizationally unique identifier (OUI)

An IEEE-standards number that identifies an organization with a 24-bit globally unique assigned number referenced by various standards. OUI is used in the family of 802 LAN standards, such as Ethernet and Token Ring.

original equipment manufacturer's information (OEMI)

A reference to an IBM guideline for a computer peripheral interface. The interface uses ESA/390 logical protocols over an I/O interface that configures attached units in a multidrop bus topology.

OS/390

The IBM operating system that includes and integrates functions that many IBM software products (including the MVS operating system) previously provided for the IBM S/390 family of enterprise servers.

OUI See *organizationally unique identifier*.

P

panel The formatted display of information that appears on a display screen.

parallel access volume (PAV)

An advanced function of the DS8000 that enables OS/390 and z/OS systems to issue concurrent I/O requests against a count key data logical volume by

associating multiple devices of a single control-unit image with a single logical device. Up to eight device addresses can be assigned to a PAV. The PAV function enables two or more concurrent write operations to the same logical volume, as long as the write operations are not to the same extents. See also *extent*, *I/O Priority Queueing*, and *multiple allegiance*.

parity A data checking scheme used in a computer system to ensure the integrity of the data. The RAID implementation uses parity to re-create data if a disk drive fails.

path group
In zSeries architecture, a set of channel paths that are defined to a control unit as being associated with a single logical partition (LPAR). The channel paths are in a group state and are online to the host. See also *logical partition*.

path group identifier
In zSeries architecture, the identifier that uniquely identifies a given logical partition (LPAR). The path group identifier is used in communication between the LPAR program and a device. The identifier associates the path group with one or more channel paths, thereby defining these paths to the control unit as being associated with the same LPAR. See also *logical partition*.

PAV See *parallel access volume*.

PCI See *peripheral component interconnect*.

PDU See *protocol data unit*.

PE See *IBM product engineering*.

peripheral component interconnect (PCI)
An architecture for a system bus and associated protocols that supports attachments of adapter cards to a system backplane.

persistent FlashCopy
A state where a FlashCopy relationship remains indefinitely until the user deletes it. The relationship between the source and target volumes is maintained after a background copy completes.

physical path
A single path through the I/O interconnection fabric that attaches two units. For Copy Services, this is the path from a host adapter on one DS8000

(through cabling and switches) to a host adapter on another DS8000.

pinned data
Data that is held in cache until either an error condition is corrected and it can be moved to disk storage or until the data is discarded by a host command. Pinned data conditions can only occur on an ESS Model 800 during fast-write or dual-copy functions.

point-in-time copy
A FlashCopy option that creates an instantaneous view of original source data at a specific moment in time.

point-to-point connection
A fibre-channel topology that enables the direct interconnection of ports. See also *arbitrated loop* and *switched fabric*.

port A physical connection on a host adapter to the cable that connects the DS8000 to hosts, switches, or another DS8000. The DS8000 uses SCSI and ESCON host adapters that have two ports per adapter, and fibre-channel host adapters that have one port. See also *ESCON*, *fibre channel*, *host adapter*, and *small computer system interface*.

POST See *power-on self test*.

power-on self test (POST)
A diagnostic test that servers or computers run when they are turned on.

predictable write
A write operation that can cache without knowledge of the existing format on the medium. All write operations on FBA DASD devices are predictable. On CKD DASD devices, a write operation is predictable if it does a format write operation for the first data record on the track.

primary control unit
The DS8000 to which a Remote Mirror and Copy primary device is physically attached.

processor complex
A partition of a storage server that is capable of performing all defined functions of the storage server. Multiple processor complexes provide redundancy.

product engineering
See *IBM product engineering*.

program

On a computer, a generic term for software that controls the operation of the computer. Typically, the program is a logical assemblage of software modules that perform multiple related tasks.

program-controlled interruption

An interruption that occurs when an I/O channel fetches a channel command word with the program-controlled interruption flag on.

program temporary fix (PTF)

A temporary solution to, or bypass of, a problem diagnosed by IBM as the result of a defect in a current unaltered release of a licensed program. (GC)

promote

To add a logical data unit to cache memory.

protected volume

In AS/400, a disk storage device that is protected from data loss by RAID techniques. An AS/400 host does not mirror a volume configured as a protected volume, while it does mirror all volumes configured as unprotected volumes. The DS8000, however, can be configured to indicate that an AS/400 volume is protected or unprotected and give it RAID protection in either case.

protocol data unit (PDU)

A unit of data specified in the protocol of a given layer and consisting of protocol control information for the layer and, possibly, user data for the layer.

pSeries

The product name of an IBM e(logo)server product that emphasizes performance. It is the successor to the RS/6000 family of servers.

pseudo host

A host connection that is not explicitly defined to the DS8000 and that has access to at least one volume that is configured on the DS8000. The FiconNet pseudo host icon represents the FICON protocol. The EsconNet pseudo host icon represents the ESCON protocol. The pseudo host icon labelled Anonymous represents hosts connected through the FCP protocol. *Anonymous host* is a commonly used synonym for *pseudo host*.

The DS8000 adds a pseudo host icon only when it is set to access-any mode. See also *access-any mode*.

PTF See *program temporary fix*.

PV Links

Short for Physical Volume Links, an alternate pathing solution from Hewlett-Packard that provides for multiple paths to a volume, as well as static load balancing.

R

R0 See *track-descriptor record*.

rack See *enclosure*.

RAID See *redundant array of independent disks*. RAID is also commonly expanded to redundant array of *inexpensive* disks. See also *array*.

RAID 5

A type of RAID that optimizes cost-effective performance while emphasizing use of available capacity through data striping. RAID 5 provides fault tolerance for up to two failed disk drives by distributing parity across all the drives in the array plus one parity disk drive. The DS8000 automatically reserves spare disk drives when it assigns arrays to a device adapter pair (DA pair). See also *device adapter*, *RAID 10*, and *redundant array of independent disks*.

RAID 10

A type of RAID that optimizes high performance while maintaining fault tolerance for up to two failed disk drives by striping volume data across several disk drives and mirroring the first set of disk drives on an identical set. The DS8000 automatically reserves spare disk drives when it assigns arrays to a device adapter pair (DA pair). See also *device adapter*, *RAID 5*, and *redundant array of independent disks*.

random access

A mode of accessing data on a medium in a manner that requires the storage device to access nonconsecutive storage locations on the medium.

rank One or more arrays that are combined to create a logically contiguous storage space.

redundant array of independent disks (RAID)

A methodology of grouping disk drives for managing disk storage to insulate data from a failing disk drive.

refresh FlashCopy target volume

An option (previously called *incremental FlashCopy*) of the DS8000 that creates a point-in-time data copy without copying an entire volume for each point-in-time copy.

Remote Mirror and Copy

A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also *Global Mirror*, *Metro Mirror* and *Global Copy*.

remote technical assistance information network (RETAIN)

The initial service tracking system for IBM service support, which captures heartbeat and call-home records. See also *support catcher* and *support catcher telephone number*.

REQ/ACK

See *request for acknowledgment and acknowledgment*.

request for acknowledgment and acknowledgment (REQ/ACK)

A cycle of communication between two data transport devices for the purpose of verifying the connection, which starts with a request for acknowledgment from one of the devices and ends with an acknowledgment from the second device. The REQ and ACK signals help to provide uniform timing to support synchronous data transfer between an initiator and a target. The objective of a synchronous data transfer method is to minimize the effect of device and cable delays.

reserved allegiance

For zSeries, a relationship that is created in a control unit between a device and a channel path, or path group, when the device completes a Sense Reserve command. The allegiance causes the control unit to guarantee access (that is, busy status is not presented) to the device. Access is over the set of channel paths that are associated with the

allegiance; access is for one or more channel programs until the allegiance ends.

RETAIN

See *remote technical assistance information network*.

S

S/390 IBM enterprise servers based on Enterprise Systems Architecture/390 (ESA/390). *S/390* is the currently accepted shortened form of the original name *System/390*.

S/390 storage

Storage arrays and logical volumes that are defined as connected to S/390 servers. This term is synonymous with count-key-data storage.

SAID See *system adapter identification number*.

SAM See *sequential access method*.

SAN See *storage area network*.

SBCON

See *Single-Byte Command Code Sets Connection*.

screen

The physical surface of a display device upon which information is shown to users.

SCSI See *small computer system interface*.

SCSI device

A disk drive connected to a host through an I/O interface using the SCSI protocol. A SCSI device is either an initiator or a target. See also *initiator* and *small computer system interface*.

SCSI-FCP

Synonym for fibre-channel protocol, a protocol used to transport data between an open-systems host and a fibre-channel adapter on an DS8000. See also *fibre-channel protocol* and *small computer system interface*.

SCSI host systems

Host systems that are attached to the DS8000 with a SCSI interface. Such host systems run on UNIX, i5/OS, Windows NT, Windows 2000, or Novell NetWare operating systems.

SCSI ID

A unique identifier assigned to a SCSI

device that is used in protocols on the SCSI interface to identify or select the device. The number of data bits on the SCSI bus determines the number of available SCSI IDs. A wide interface has 16 bits, with 16 possible IDs.

SDD See *IBM Subsystem Multipathing Device Driver*.

secondary control unit

The DS8000 to which a Remote Mirror and Copy secondary device is physically attached.

self-timed interface (STI)

An interface that has one or more conductors that transmit information serially between two interconnected units without requiring any clock signals to recover the data. The interface performs clock recovery independently on each serial data stream and uses information in the data stream to determine character boundaries and inter-conductor synchronization.

sequential access

A mode of accessing data on a medium in a manner that requires the storage device to access consecutive storage locations on the medium.

sequential access method (SAM)

An access method for storing, deleting, or retrieving data in a continuous sequence based on the logical order of the records in the file.

serial connection

A method of device interconnection for determining interrupt priority by connecting the interrupt sources serially.

server A host that provides certain services to other hosts that are referred to as clients.

A functional unit that provides services to one or more clients over a network. (GC)

service boundary

A category that identifies a group of components that are unavailable for use when one of the components of the group is being serviced. Service boundaries are provided on the DS8000, for example, in each host bay and in each cluster.

service clearance

The area that is required to open the service covers and to pull out components for servicing.

service information message (SIM)

A message sent by a storage server to service personnel through an zSeries operating system.

service personnel

A generalization referring to individuals or companies authorized to service the DS8000. The terms *service provider*, *service representative*, and *IBM service support representative (SSR)* refer to types of service personnel. See also *service support representative*.

service processor

A dedicated processing unit that is used to service a storage unit.

service support representative (SSR)

Individuals or a company authorized to service the DS8000. This term also refers to a service provider, a service representative, or an IBM service support representative (SSR). An IBM SSR installs the DS8000.

SES SCSI Enclosure Services.

session

A collection of volumes within a logical subsystem that are managed together during the creation of consistent copies of data. All volumes in a session must transfer their data successfully to the remote site before the increment can be called complete.

SFP Small form factor pluggables.

shared storage

Storage that is configured so that multiple hosts can concurrently access the storage. The storage has a uniform appearance to all hosts. The host programs that access the storage must have a common model for the information on a storage device. The programs must be designed to handle the effects of concurrent access.

shortwave laser adapter

A connector that is used between host and DS8000 to support shortwave fibre-channel communication.

SIM See *service information message*.

Simple Network Management Protocol (SNMP)

In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application layer protocol. Information on devices managed is defined and stored in the application's Management Information Base (MIB). (GC) See also *management information base*.

simplex volume

A volume that is not part of a FlashCopy, XRC, or PPRC volume pair.

Single-Byte Command Code Sets Connection (SBCON)

The ANSI standard for the ESCON I/O interface.

small computer system interface (SCSI)

A standard hardware interface that enables a variety of peripheral devices to communicate with one another. (GC)

smart relay host

A mail relay or mail gateway that has the capability to correct e-mail addressing problems.

SMIT See *System Management Interface Tool*.

SMP See *symmetrical multiprocessor*.

SNMP See *Simple Network Management Protocol*.

SNMP agent

A server process that resides on a network node and is responsible for communicating with managers regarding that node. The node is represented as a managed object, which has various fields or variables that are defined in the appropriate MIB.

SNMP manager

A managing system that runs a managing application or suite of applications. These applications depend on Management Information Base (MIB) objects for information that resides on the managed system. Managers generate requests for this MIB information, and an SNMP agent on the managed system responds to these requests. A request can either be the retrieval or modification of MIB information.

software transparency

Criteria applied to a processing

environment that states that changes do not require modifications to the host software in order to continue to provide an existing function.

source device

One of the devices in a dual-copy or remote-copy volume pair. All channel commands to the logical volume are directed to the source device. The data on the source device is duplicated on the target device. See also *target device*.

spare A disk drive on the DS8000 that can replace a failed disk drive. A spare can be predesignated to allow automatic dynamic sparing. Any data preexisting on a disk drive that is invoked as a spare is destroyed by the dynamic sparing copy process.

spatial reuse

A feature of serial storage architecture that enables a device adapter loop to support many simultaneous read/write operations. See also *serial storage architecture*.

SSID See *subsystem identifier*.

SSR See *service support representative*.

stacked status

For zSeries, the condition when the control unit is in a holding status for the channel, and the last time the control unit attempted to present the status, the channel responded with the stack-status control.

stage operation

The operation of reading data from the physical disk drive into the cache.

staging

To move data from an offline or low-priority device back to an online or higher priority device, usually on demand of the system or on request of the user.

standard volume

A volume that emulates one of several zSeries volume types, including 3390-2, 3390-3, 3390-9, 3390-2 (3380-track mode), or 3390-3 (3380-track mode), by presenting the same number of cylinders and capacity to the host as provided by the native zSeries volume type of the same name.

STI See *self-timed interface*.

storage area network

A network that connects a company's heterogeneous storage resources.

storage capacity

The amount of data that a storage medium can hold; usually expressed in kilobytes, megabytes, or gigabytes.

storage complex

A configuration of one or more storage units that is managed by a management console.

storage device

A physical unit that provides a mechanism to store data on a given medium such that it can be subsequently retrieved. See also *disk drive module*.

storage extent

The minimum contiguous range of storage on a physical storage device, array, or rank that can be allocated to a local volume

storage image

A partitioning of a storage unit that provides emulation of a storage server with one or most storage devices that provides storage capability to a host computer. You can configure more than one storage image on a storage unit.

storage server

A physical unit that manages attached storage devices and provides an interface between them and a host computer by providing the function of one or more logical subsystems. The storage server can provide functions that the storage device does not provide. The storage server has one or more clusters.

storage unit

A physical unit that consists of a storage server that is integrated with one or more storage devices that provide storage capability to a host computer.

storage unit identifier

A unique identifier for a storage unit that consists of a manufacturer, a model number, a type number, a plant of manufacture, and a sequence number.

striping

A technique that distributes data in bit, byte, multibyte, record, or block increments across multiple disk drives.

subagent

An extension to an SNMP agent that permits a user to dynamically add, or in some cases replace, additional management variables in the local MIB, thereby providing a means of extending the range of information that network managers can access. See also *agent*.

subchannel

A logical function of a channel subsystem associated with the management of a single device.

subordinate storage unit

The physical unit that receives commands from the master storage unit and is specified when a Global Mirror session is started. The subordinate storage unit forms consistency groups and performs other Global Mirror processing. A subordinate storage unit can be controlled by only one master storage unit. Contrast with *master storage unit*.

subsystem identifier (SSID)

A number that uniquely identifies a logical subsystem within a computer installation.

support catcher

See *catcher*.

support catcher telephone number

The telephone number that connects the support catcher server to the DS8000 to receive a trace or dump package. See also *support catcher* and *remote technical assistance information network*.

switched fabric

A fibre-channel topology in which ports are interconnected through a switch. Fabric switches can also be interconnected to support numerous ports on a single network. See also *arbitrated loop* and *point-to-point connection*.

symmetrical multiprocessor (SMP)

An implementation of a multiprocessor computer consisting of several identical processors configured in a way that any subset of the set of processors is capable of continuing the operation of the computer. The DS8000 contains four processors set up in SMP mode.

synchronous write

A write operation whose completion is indicated after the data has been stored on a storage device.

System/390

See *S/390*.

system adapter identification number (SAID)

The unique identification number that is automatically assigned to each DS8000 host adapter for use by Copy Services.

System Management Interface Tool (SMIT)

An interface tool of the AIX operating system for installing, maintaining, configuring, and diagnosing tasks.

System Modification Program

A program used to install software and software changes on MVS systems.

T

TAP See *Telocator Alphanumeric Protocol*.

target A SCSI device that acts as a subordinate to an initiator and consists of a set of one or more logical units, each with an assigned logical unit number (LUN). The logical units on the target are typically I/O devices. A SCSI target is analogous to a zSeries control unit. See also *small computer system interface*.

target device

One of the devices in a dual-copy or remote-copy volume pair that contains a duplicate of the data that is on the source device. Unlike the source device, the target device might only accept a limited subset of data. See also *source device*.

TB See *terabyte*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

Telocator Alphanumeric Protocol (TAP)

An industry standard protocol for the input of paging requests.

terabyte (TB)

- 1) Nominally, 1 000 000 000 000 bytes, which is accurate when speaking of bandwidth and disk storage capacity.
- 2) For DS8000 cache memory, processor storage, real and virtual storage, a terabyte refers to 2⁴⁰ or 1 099 511 627 776 bytes.

terminal emulator

A function of the management console that allows it to emulate a terminal.

thousands of power-on hours (KPOH)

A unit of time used to measure the mean time between failures (MTBF).

time sharing option (TSO)

An operating system option that provides interactive time sharing from remote terminals.

TotalStorage

See *IBM TotalStorage*.

TPF See *transaction processing facility*.

track A unit of storage on a CKD device that can be formatted to contain a number of data records. See also *home address*, *track-descriptor record*, and *data record*.

track-descriptor record (R0)

A special record on a track that follows the home address. The control program uses it to maintain certain information about the track. The record has a count field with a key length of zero, a data length of 8, and a record number of 0. This record is sometimes referred to as R0.

transaction processing facility (TPF)

A high-availability, high-performance IBM operating system, designed to support real-time, transaction-driven applications. The specialized architecture of TPF is intended to optimize system efficiency, reliability, and responsiveness for data communication and database processing. TPF provides real-time inquiry and updates to a large, centralized database, where message length is relatively short in both directions, and response time is generally less than three seconds. Formerly known as the Airline Control Program/Transaction Processing Facility (ACP/TPF).

Transmission Control Protocol (TCP)

A communications protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

Transmission Control Protocol/Internet Protocol (TCP/IP)

- 1) A combination of data-transmission

protocols that provide end-to-end connections between applications over interconnected networks of different types.

2) A suite of transport and application protocols that run over the Internet Protocol. (GC) See also *Internet Protocol* and *Transmission Control Protocol*.

transparency

See *software transparency*.

TSO See *time sharing option*.

turbo processor

A faster multiprocessor that has six processors with common access to the main storage.

U

UFS UNIX filing system.

Ultra-SCSI

An enhanced small computer system interface.

unconfigure

To delete the configuration.

unit address

For zSeries, the address associated with a device on a given control unit. On ESCON interfaces, the unit address is the same as the device address. On OEMI interfaces, the unit address specifies a control unit and device pair on the interface.

unprotected volume

An AS/400 term that indicates that the AS/400 host recognizes the volume as an unprotected device, even though the storage resides on a RAID-formatted array and is, therefore, fault tolerant by definition. The data in an unprotected volume can be mirrored. Also referred to as an *unprotected device*.

upper-layer protocol

The layer of the Internet Protocol (IP) that supports one or more logical protocols (for example, a SCSI-command protocol and an ESA/390 command protocol). Refer to ANSI X3.230-199x.

UTC See *Coordinated Universal Time*.

V

virtual machine facility

A virtual data processing machine that appears to the user to be for the exclusive

use of that user, but whose functions are accomplished by sharing the resources of a shared data processing system. An alternate name for the VM/370 IBM operating system.

vital product data (VPD)

Information that uniquely defines the system, hardware, software, and microcode elements of a processing system.

VM

The root name of several IBM operating systems, such as VM/XA, VM/ESA, VM/CMS, and z/VM. See also *virtual machine facility*.

volume

For zSeries, the information recorded on a single unit of recording medium. Indirectly, it can refer to the unit of recording medium itself. On a nonremovable-medium storage device, the term can also indirectly refer to the storage device associated with the volume. When multiple volumes are stored on a single storage medium transparently to the program, the volumes can be referred to as logical volumes.

volume group

A collection of either physical or logical volumes.

volume label

A unique identifier that a user assigns to a logical volume.

VPD See *vital product data*.

VSE/ESA

An IBM operating system, the letters of which represent virtual storage extended/enterprise systems architecture.

W

weight distribution area

The area that is required to distribute the weight of the storage unit.

worldwide node name (WWNN)

A unique 64-bit identifier for a host that contains a fibre-channel port. See also *worldwide port name*.

worldwide port name (WWPN)

A unique 64-bit identifier associated with a fibre-channel adapter port. It is assigned in an implementation- and

protocol-independent manner. See also *worldwide node name*

write hit

A write operation in which the requested data is in the cache.

write penalty

The performance impact of a classical RAID-5 write operation.

WWNN

See *worldwide node name*.

WWPN

See *worldwide port name*.

X

xSeries

The product name of an IBM e(logo)server product that emphasizes industry-standard server scalability and self-managing server technologies. It is the successor to the Netfinity family of servers.

Z

z/Architecture

An IBM architecture for mainframe computers and peripherals. The IBM e(logo)server zSeries family of servers uses the z/Architecture architecture. It is the successor to the S/390 and 9672 family of servers. See also *iSeries*.

z/OS

An operating system for the IBM e(logo)server product line that supports 64-bit real storage.

z/OS Global Mirror

A function of a storage server that assists a control program to maintain a consistent copy of a logical volume on another storage unit. All modifications of the primary logical volume by any attached host are presented in order to a single host. The host then makes these modifications on the secondary logical volume. This function was formerly called extended remote copy or XRC.

zSeries

An IBM e(logo)server family of servers that emphasizes near-zero downtime.

IBM enterprise servers based on z/Architecture.

zSeries storage

Storage arrays and logical volumes that are defined in the DS8000 as connected to zSeries servers.

Index

Numerics

- 2000 host system, Windows
 - configuring for availability and recoverability 206
- 2003 host system, configuring Windows Server for
 - availability and recoverability 220

A

- access profiles 21
- accessibility features 233
 - keyboard 233
 - shortcut keys 233
- adapter card
 - fibre-channel 19
- adapter cards
 - 2766 IBM iSeries host system 70
 - driver, downloading 205, 219
 - ESCON host 100
 - Hewlett-Packard host system
 - A5158A 65
 - A6684A 65
 - A6685A 65
 - A6795A 65
 - A6826A 65
 - A9782A 65
 - installing PCI-FC-1P-OPT 166
 - installing PCI-FC-1P-OPT-A 166
 - Intel host system
 - Emulex LP10000 119
 - Emulex LP10000DC 119
 - Emulex LP9002L 119
 - Emulex LP9402DC 119
 - Emulex LP9802 119
 - QLogic QLA2310FL 119
 - QLogic QLA2340FL 119
 - QLogic QLA2342FL 119
 - KGPSA-CA
 - Hewlett-Packard host system 53
 - HP host system 40
 - KGPSA-DA
 - Hewlett-Packard host system 53
 - HP host system 40
 - KGPSA-EA
 - Hewlett-Packard host system 53
 - HP host system 40
 - Linux host system
 - Emulex LP10000 119
 - Emulex LP10000DC 119
 - Emulex LP9002L 119
 - Emulex LP9402DC 119
 - Emulex LP9802 119
 - QLogic QLA2310FL 119
 - QLogic QLA2340FL 119
 - QLogic QLA2342FL 119
 - SGL host system
 - PCI-FC-1P-OPT 165
 - PCI-FC-1P-OPT-A 165
 - adapter cards (*continued*)
 - SGL host system (*continued*)
 - PCI-FC-1P-OPT-B 165
 - PCX-FC-2P-OPT-B 165
 - Windows Server 2003 host system
 - Emulex LP10000 211
 - Emulex LP10000DC 211
 - Emulex LP9002L 211
 - Emulex LP9402DC 211
 - Emulex LP9802 211
 - Netfinity 211
 - QLogic QLA2310FL 211
 - QLogic QLA2340FL 211
 - QLogic QLA2342FL 211
 - adapter cards for IBM NAS Gateway 500 hosts,
 - Supported 79
 - adapter driver
 - downloading QLogic 123
 - HP-UX11.i, and HP-UX 11iv2 installing 66
 - installing 32, 181
 - installing Emulex 199, 213
 - installing the AMCC PCI fibre channel 182
 - installing the AMCC SBUS fibre channel 182
 - adapter drivers, installing Novell NetWare 154
 - adapter ports
 - configuring for Hewlett-Packard 56
 - configuring for HP 42
 - adapters
 - attaching to open-systems host 13
 - adapters, general information about attaching to an open-systems hosts 16
 - adding
 - AlphaServer connections
 - Hewlett-Packard host system 55
 - HP host system 42
 - adding volumes
 - Tru64 UNIX 42
 - affinity for LUNs 20
 - AIX 5L Version 5.2, Migrating with 96
 - AIX hosts, Upgrading the host attachment package on IBM pSeries 86
 - AlphaServer connections
 - adding 42, 55
 - defining 55
 - modifying 42, 55
 - AlphaServer console, support 38, 39
 - AMCC PCI adapter card
 - downloading, Sun host system 182
 - installing, Sun host system 181
 - AMCC PCI fibre-channel adapter driver, installing 182
 - AMCC SBUS adapter card
 - downloading, Sun host system 182
 - installing, Sun host system 182
 - AMCC SBUS fibre-channel adapter driver,
 - installing 182
 - anonymous host 21
 - API, DS Open 9

- arbitrated loop topology
 - description 16
 - illustration of 16
 - assigning volumes
 - Tru64 UNIX 42
 - attaching
 - FICON channel 109
 - FICON channel on a S/390 or zSeries host 110
 - Hewlett-Packard host system 52
 - Hewlett-Packard OpenVMS 52
 - HP host system 40
 - HP Tru64 UNIX Version 5.0x 40
 - Intel host system 119, 139
 - iSCSI Gateway host system 157
 - iSeries host system 69
 - Linux host system 119
 - multiple RS/6000 or pSeries hosts without HACMP 91
 - Novell NetWare host system 151
 - pSeries host system 85, 91
 - RS/6000 host system 85, 91
 - S/390 host system 99
 - SAN File System 161
 - SAN Volume Controller host system 163
 - SGI host system 165
 - to an open-systems hosts with fibre-channel adapters 16
 - VMware ESX server host system 139
 - Windows 2000 host system 197
 - Windows Server 2003 host system 211
 - with ESCON adapters
 - S/390 100
 - zSeries 100
 - with fibre-channel adapters
 - Hewlett-Packard 51, 65
 - HP 37
 - Intel 119, 139
 - iSeries 69
 - Linux 119
 - Novell NetWare 151
 - PRIMEPOWER 31
 - RS/6000 86
 - SAN Volume Controller 163
 - SGI 165
 - Sun 179
 - VMware ESX server 139
 - Windows 2000 197
 - Windows Server 2003 211
 - xSeries 86
 - with fibre—channel adapters
 - Hewlett-Packard host system 65
 - zSeries host system 99
 - Attaching to an IBM NAS Gateway 500 Host 79
 - attachment
 - configuring the storage unit for FICON 108
 - considerations
 - FICON 108
 - iSCSI Gateway router 158
 - overview for iSCSI router 157
 - requirements
 - iSCSI router 157
 - attachment (*continued*)
 - requirements (*continued*)
 - SGI 165
 - RS/6000 and pSeries AIX host system, installing package 85
 - storage unit volumes for Tru64 UNIX 45
 - attachment considerations
 - Hewlett-Packard host system
 - fibre-channel 51
 - HP host system
 - fibre-channel 38
 - Intel host system, fibre-channel 120, 140
 - iSeries host system, fibre-channel 69
 - Linux host system, fibre-channel 120
 - Novell NetWare host system, fibre-channel 151
 - pSeries host system, fibre-channel 87
 - RS/6000 host system
 - fibre-channel 87
 - SGI host system fibre-channel 165
 - VMware ESX Server host system,
 - fibre-channel 140
 - Windows 2000 host system, fibre-channel 198
 - Windows 2003 host system, fibre-channel 212
 - attachment multipathing scripts, Host 81
 - attachment requirements
 - Hewlett-Packard host system
 - fibre-channel 51, 65
 - HP host system
 - fibre-channel 37
 - Intel host system, fibre-channel 119, 139
 - iSeries host system, fibre-channel 69
 - Linux host system, fibre-channel 119
 - PRIMEPOWER host system, fibre-channel 31
 - pSeries host system
 - fibre-channel 87
 - RS/6000 host system
 - fibre-channel 87
 - Sun host system
 - fibre-channel 180
 - VMware ESX Server host system,
 - fibre-channel 139
 - Windows 2000 host system
 - fibre-channel 197
 - Windows Server 2003 host system
 - fibre-channel 211
 - attachment, multipath 70, 73
 - availability
 - configuring for 206, 220
 - availability features 11
- ## B
- boot considerations for IBM pSeries hosts,
 - Fibre-channel 89
 - boot support for a Windows 2000 host system, Installing remote fibre-channel 207
 - boot support for a Windows Server 2003 host system, Installing remote fibre-channel 221
 - boot, Support for fibre-channel 130
 - bridge configuration, FICON 105

C

- cable
 - distances for S/390 and zSeries, ESCON 100
 - ESCON host 100
 - fibre-channel 19
 - installing SGI 167
 - lengths
 - for S/390 and zSeries 101
 - SGI installing 167
- changing the Sun system kernel 184
- channel
 - director 111
 - directors for S/390 and zSeries 101
 - extenders for S/390 and zSeries 101
- checking
 - installation of PCI-FC-1P-OPT adapter card 166
 - installation of PCI-FC-1P-OPT-A adapter card 166
 - level of the device driver for Netfinity adapter card 202, 216
 - version of the IRIX operating system, SGI 166
- command-line interface 9
- command, port identification for S/390 and zSeries TSO 101
- comparing
 - Copy Services
 - features 27
- configuration
 - FICON 104, 107
 - FICON bridge 105
 - iSeries 103
 - native ESCON 102
 - S/390 103
- configurations for RS/6000 for a fibre-channel boot install, Support for disk 90
- configuring
 - drivers for SGI 167
 - for availability and recoverability for a Windows 2000 operating system 206
 - for availability and recoverability for a Windows Server 2003 operating system 220
 - host device drivers for PRIMEPOWER 32
 - host device drivers for Sun 184
 - IBM iSeries 72
 - MC/ServiceGuard on an HP-UX with the storage unit 66
 - storage unit for FICON attachment 108
 - storage unit for host failover, SGI 170, 171
 - storage unit on a Linux 125
 - Windows 2000
 - configuration settings for Emulex adapter cards for 200
 - Windows Server 2003
 - configuration settings for Emulex adapter cards for 214
- Configuring 161
- configuring host adapter ports
 - Hewlett-Packard 56
 - HP 42
- configuring storage 46, 61
 - storage for SGI 172

- confirming
 - storage connectivity, SGI 169, 170
 - switch connectivity, SGI 167
- connecting
 - fibre-channel 16
- considerations
 - for fibre-channel attachment
 - Hewlett-Packard 51
 - HP 38
 - Intel 120, 140
 - iSeries 69
 - Linux 120
 - Novell NetWare 151
 - pSeries 87
 - RS/6000 87
 - SGI 165
 - VMware ESX Server 140
 - Windows 2000 198
 - Windows 2003 212
- considerations for IBM pSeries hosts, Fibre-channel boot 89
- considerations for NAS Gateway 500, Multipath I/O and SDD 81
- console, support for the AlphaServer 38, 39
- controller
 - images for S/390 and zSeries 100
 - S/390 and zSeries 100
- Copy Services
 - disaster recovery 26
 - feature comparison 27
 - functions of 23
 - monitoring 23
 - overview 23
- creating file systems on the storage unit, Linux 127

D

- data
 - restoring 92
 - saving 91
 - transferring 101
- defining
 - AlphaServer OpenVMS connections, HP host system 55
 - number of disk devices on Linux 124
- Device Driver, Installing Linux over the SAN without an IBM Subsystem 131
- device drivers, configuring 32, 184
- director, channel 111
- directors and channel extenders for S/390 and zSeries 101
- disaster recovery
 - using Copy Services 26
- disk configurations for RS/6000 for a fibre-channel boot install, Support for 90
- disk devices on Linux, defining 124
- disks, partitioning storage units for Linux 125
- displaying the storage unit volume 43, 59
- displaying zoning information for the switch, SGI 168
- distances, fibre-channel node-to-node 19

- downloading
 - AMCC PCI fibre-channel adapter driver 182
 - AMCC SBUS fibre-channel adapter driver 182
 - Emulex adapter driver 32, 152, 181, 199, 213
 - Emulex fibre-channel adapter driver 121
 - QLogic adapter driver 152, 183
- downloading the current fibre-channel adapter driver 123, 205, 219
- drivers
 - configuring SGI 167
 - installing
 - AMCC PCI fibre-channel adapter 182
 - AMCC SBUS fibre-channel adapter 182
 - installing for Novell NetWare 154
 - installing QLogic 123
 - Windows 2000, installing 205
 - Windows Server 2003, installing 219
- drivers, configuring SGI 167
- DS command-line interface 9
- DS Open API 9
- DS Storage Manager 9
- DS8000 series
 - Model 921 1
 - Model 922 1
 - Model 92E 1
 - Model 9A2 1
 - Model 9AE 1
- DS8100 1
- DS8300 1

E

- Emulex
 - adapter settings
 - configuring for Windows 2000 200
 - configuring for Windows Server 2003 214
- Emulex adapter card
 - downloading 32, 152, 181, 199, 213
 - installing
 - Linux host system 121
 - Novell NetWare host system 151
 - PRIMEPOWER host system 31
 - Sun host system 180
 - Windows 2000 host system 198, 199
 - Windows Server 2003 host system 213
- Emulex adapter driver
 - downloading 121
 - installing 32, 152, 181
- Emulex adapter drivers, installing 121, 140, 199, 213
- ESCON
 - cable distances for S/390 and zSeries host system 100
 - controller images and controls 100
 - host adapters for 100
 - host cables 100
 - host systems 15
 - multipathing 109
- ESCON adapters, attaching with
 - S/390 100
 - zSeries 100
- ESCON configuration, native 102

- ESCON distances 100
- ESCON links, setting up 108
- extenders for S/390 and zSeries
 - channels 101
 - directors 101

F

- failover, configuring the storage unit 170
- failover, configuring the storage unit for SGI 171
- FCP for a zSeries host, Accessing the storage unit using 114
- fibre channel
 - adapters 13
- fibre-channel
 - adapter 16
 - adapter driver, downloading 123, 205, 219
 - adapter drivers
 - installing the AMCC PCI 182
 - installing the AMCC SBUS 182
 - adapters 19
 - arbitrated loop 16
 - cable 16
 - cables 19
 - connection 16
 - host systems 14
 - intermix mode 111
 - loop initialization 16
 - LUN access modes 20
 - LUN affinity 20
 - node-to-node distances 19
 - port name identification 225
 - ports 16
 - storage area networks (SANs) 21
 - targets and LUNs 20
 - topologies 16
- fibre-channel access modes 20
- Fibre-channel adapters, Attaching to an IBM SAN File System metadata server 161
- fibre-channel adapters, attaching with
 - Hewlett-Packard 51, 65
 - Hewlett-Packard host system 65
 - HP 37
 - Intel 119, 139
 - iSeries 69
 - Linux 119
 - Novell NetWare 151
 - PRIMEPOWER 31
 - RS/6000 86
 - SAN Volume Controller 163
 - SGI 165
 - Sun 179
 - VMware ESX server 139
 - Windows 2000 197
 - Windows Server 2003 211
 - xSeries 86
- fibre-channel adapters, general information about attaching 16
- fibre-channel arbitrated loop, SGI 170, 172, 175
- fibre-channel boot install, Support for disk configurations for RS/6000 for a 90

G

H

- Host attachment multipathing scripts 81
- host attachment package on IBM pSeries AIX hosts, Upgrading the 86
- host failover when configuring the storage unit 170
- host failover when configuring the storage unit, SGI 171
- host limitations, IBM iSeries 70
- host system, Installing remote fibre-channel boot support for a Windows 2000 207
- host systems
 - attaching the storage unit
 - Intel 119, 139
 - iSCSI Gateway 157
 - iSeries 69
 - Linux 119
 - Novell NetWare 151
 - pSeries 85
 - RS/6000 85
 - S/390 99
 - SAN File System 161
 - SAN Volume Controller 163
 - SGI 165
 - VMware ESX server 139
 - Windows 2000 197
 - Windows Server 2003 211
 - zSeries 99
 - fibre channel 14
 - S/390 and zSeries 15
 - supported by the storage unit 13
 - without HACMP, attaching a storage unit to multiple RS/6000 or pSeries 91
- host, anonymous 21
- host, Attaching to an IBM NAS Gateway 500 79
- host, Running the Linux operating system on an IBM i5 74
- host, Running the Linux operating system on an IBM pSeries 92
- hosts with fibre-channel adapters, general information 16
- hosts, Fibre-channel boot considerations for IBM pSeries 89
- hosts, Registered state-change notifications (RSCNs) on zSeries 111
- hosts, Supported adapter cards for IBM NAS Gateway 500 79
- HP AlphaServer connections
 - adding 42
 - modifying 42
- HP AlphaServer OpenVMS
 - microcode levels 52
- HP AlphaServer OpenVMS connections
 - defining 55
- HP host system
 - attachment considerations
 - fibre-channel 38
 - attachment requirements
 - fibre-channel 37
 - Tru64 UNIX
 - Version 5.0x 40
- HP OpenVMS, microcode levels 52
- HP-UX 11i, installing adapter drivers 66

HP-UX 11iv2, installing adapter drivers 66

I

i5/OS Version 5 Release 3 69, 70, 73

IBM i5 server, Running the Linux operating system on an 74

IBM iSeries

hardware 71

software 71

IBM iSeries host system, recommended configurations 72

IBM NAS Gateway 500 host, Attaching to an 79

IBM NAS Gateway 500 hosts, Supported adapter cards for 79

IBM pSeries host, Running the Linux operating system on an 92

IBM Subsystem Device Driver, Installing Linux over the SAN with an 132

IBM Subsystem Device Driver, Installing Linux over the SAN without an 131

image for S/390 and zSeries 100

install, Support for disk configurations for RS/6000 for a fibre-channel boot 90

installation

preparing for RS/6000 and pSeries AIX hosts 85

installation package for RS/6000 and pSeries AIX hosts 85

installing

adapter drivers for HP-UX 11.i, and HP11iv2 66

adapter drivers, Novell NetWare 154

AMCC PCI adapter card 181

AMCC PCI adapter driver, fibre-channel 182

AMCC SBUS adapter card 182

AMCC SBUS adapter driver, fibre-channel 182

device driver for a Netfinity adapter card 203, 217

drivers for Windows 2000 205

drivers for Windows Server 2003 219

Emulex adapter card 31, 151, 180, 198, 199, 213

Emulex adapter card on a Linux host system 121

Emulex adapter driver 32, 152, 181

Emulex adapter drivers 121, 140, 199, 213

fibre-channel adapter drivers 205, 219

fibre-channel adapter drivers for SGI 167

host attachment package for an RS/6000 and pSeries AIX hosts 85

host attachment package for RS/6000 and pSeries AIX host systems 85

KGPSA-CA adapter card

Hewlett-Packard host system 53

HP host system 40

KGPSA-DA adapter card

Hewlett-Packard host system 53

HP host system 40

KGPSA-EA adapter card

Hewlett-Packard host system 53

HP host system 40

Netfinity adapter card 201

optical cable, SGI 167

optical SGI cable 167

QLogic adapter card 153, 183

installing (*continued*)

QLogic adapter card on a Linux host system 122

QLogic adapter card on a VMware ESX Server host system 141

QLogic adapter drivers 123

QLogic QLA23xx adapter card 203, 217

SDD 194

SGI PCI-FC-1P-OPT 166

SGI PCI-FC-1P-OPT-A 166

Installing QLogic adapter driver package 183

Installing remote fibre-channel boot support for a Windows 2000 host system 207

Installing remote fibre-channel boot support for a Windows Server 2003 host system 221

Intel host system

attaching the storage unit 119, 139

attachment considerations, fibre-channel 120, 140

attachment requirements, fibre-channel 119, 139

Emulex LP10000 adapter card 119

Emulex LP10000DC adapter card 119

Emulex LP9002L adapter card 119

Emulex LP9402DC adapter card 119

Emulex LP9802 adapter card 119

QLogic QLA2310FL adapter card 119

QLogic QLA2340FL adapter card 119

QLogic QLA2342FL adapter card 119

introduction 1

IRIX operating system, checking the version of the SGI 166

iSCSI Gateway host system, attaching the storage unit 157

iSeries host system

attaching the storage unit 69

attachment considerations, fibre-channel 69

attachment requirements, fibre-channel 69

locating the WWPN 227

K

KGPSA-CA adapter card, installing

Hewlett-Packard host system 53

HP host system 40

KGPSA-DA adapter card, installing

Hewlett-Packard host system 53

HP host system 40

KGPSA-EA adapter card, installing

Hewlett-Packard host system 53

HP host system 40

L

levels, microcode 38

levels, operating system 37

limitations

IBM iSeries 70

limitations for Tru64 UNIX 50

links, setting up ESCON and FICON 108

Linux

creating file systems on the storage unit 127

defining the number of disk devices on 124

Red Hat 120

- Linux *(continued)*
 - SUSE 120
 - using file systems on the storage unit 127
- Linux host system
 - attaching the storage unit 119
 - attachment considerations, fibre-channel 120
 - attachment requirements, fibre-channel 119
 - Emulex LP10000 adapter card 119
 - Emulex LP10000DC adapter card 119
 - Emulex LP9002L adapter card 119
 - Emulex LP9402DC adapter card 119
 - Emulex LP9802 adapter card 119
 - QLogic QLA2310FL adapter card 119
 - QLogic QLA2340FL adapter card 119
 - QLogic QLA2342FL adapter card 119
- Linux host, locating the WWPN 228
- Linux in a guest partition, Planning to run 75
- Linux operating system on an IBM i5 server, Running the 74
- Linux operating system on an IBM pSeries host, Running the 92
- Linux over the SAN without an IBM Subsystem Device Driver, Installing 131
- Linux, Attaching a storage unit to an S/390 or zSeries host running 112
- locating the WWPN
 - Hewlett-Packard host 226
 - iSeries host 227
 - Linux host 228
 - Novell NetWare host 229
 - pSeries host 227
 - RS/6000 host 227
 - SGI host 229
 - Sun host 230
 - Windows 2000 host 230
 - Windows Server 2003 host 231
- logical paths for S/390 and zSeries 101
- loop, arbitrated 16
- LUN
 - affinity for fibre-channel 20

M

- management console
 - multiple units 12
 - overview 12
- MC/ServiceGuard on an Hewlett-Packard host system, configuring 66
- microcode levels 38
- microcode levels, HP OpenVMS 52
- migrating
 - ESCON to native FICON 102
 - FICON bridge to native FICON 105
- Migrating with AIX 5L Version 5.2 96
- MIH (missing-interrupt handler) setting 15
- missing-interrupt handler (MIH) setting 15
- mixed configuration
 - ESCON 103
 - FICON 106
- modes, fibre-channel access 20

- modifying AlphaServer connections
 - Hewlett-Packard host system 55
 - HP host system 42
- multipath attachment 70, 73
- Multipath I/O and SDD considerations for NAS Gateway 500 81
- multipath requirement 69
- multipathing
 - ESCON 109
 - FICON 109
- multipathing scripts, Host attachment 81
- Multipathing support for NAS Gateway 500 81
- Multipathing with the Subsystem Device Driver 83
- Multiple Device Manager (MDM), Copy Services 23

N

- name, Finding the worldwide port 79
- NAS Gateway 500 host, Attaching to an IBM 79
- NAS Gateway 500, Multipath I/O and SDD considerations for 81
- NAS Gateway 500, Multipathing support for 81
- native ESCON configuration 102
- native FICON configuration 107
- Netfinity adapter card
 - checking the level of the device driver 202, 216
 - installing 201
 - installing the device driver 203, 217
- node-to-node distances 19
- node-to-node distances, fibre-channel 19
- Novell NetWare host system
 - attaching the storage unit 151
 - attachment considerations, fibre-channel 151
- Novell NetWare host, locating the WWPN 229

O

- Obtaining WWPNS through the command-line interface 79
- Obtaining WWPNS using a Web browser 79
- on IBM pSeries AIX hosts, Upgrading the host attachment package 86
- open-systems
 - hosts
 - fibre-channel 14
- OpenVMS
 - attachment considerations 38, 51
 - microcode levels 52
 - storage connectivity 43, 59
 - switch connectivity 58
 - Version 7.3 52
 - Version 7.3-1 52
- operating system
 - checking the version of the IRIX, SGI 166
- operating system levels 37
- operating system on an IBM i5 server, Running the Linux 74
- operating system on an IBM pSeries host, Running the Linux 92
- optical cable, installing SGI 167

overview

Copy Services

monitor 23

Multiple Device Manager (MDM) 23

FlashCopy 25

P

package

RS/6000 and pSeries AIX hosts, installing 85

parameters

recommended configuration file parameters for
Emulex LP10000 185

recommended configuration file parameters for
Emulex LP10000DC 185

recommended configuration file parameters for
Emulex LP9002DC 185

recommended configuration file parameters for
Emulex LP9002L 185

recommended configuration file parameters for
Emulex LP9002S 185

settings for AMCC FCE-1473 187

settings for AMCC FCE-6460 187

settings for AMCC FCX-6562 187

settings for AMCC FCX2-6562 187

settings for Emulex LP9002L 34

settings for QLogic QLA2310F 190

settings for QLogic QLA2340 190

partitioning storage unit disks, Linux 125

paths

logical for S/390 and zSeries 101

types for S/390 and zSeries 101

PCI fibre-channel adapter driver, installing the

AMCC 182

PCI-FC-1P-OPT adapter card for SGI, installing 166

PCI-FC-1P-OPT adapter card, installing SGI 166

persistent reserves, removing Tru64 UNIX 47

physical configuration

I/O enclosures 1

physical configuration

battery assemblies 1

device adapters 1

disk enclosures 1

host adapters 1

management consoles 1

power line disturbance (PLD) 1

point-to-point topology

description 16

illustration of 16

port identification for S/390 and zSeries TSO

commands 101

port name identification for fibre-channel 225

ports

configuring host adapter for Hewlett-Packard 56

configuring host adapter for HP 42

primary path, switching I/O 172

PRIMEPOWER host system attachment requirements,

fibre-channel 31

processors 1

profiles, access 21

pSeries AIX hosts, Upgrading the host attachment
package 86

pSeries host system

attaching the storage unit 85

attachment requirements

fibre-channel 87

fibre-channel attachment considerations 87

pSeries host, locating the WWPN 227

pSeries hosts, Fibre-channel boot considerations for
IBM 89

Q

QLA23xx adapter card, installing 203, 217

QLogic adapter card, installing 122, 141, 153, 183

QLogic adapter driver package, installing 183

QLogic adapter drivers, downloading 152

QLogic QLA23xx adapter card, installing 203, 217

queue depth for the HP-UX operating system,
setting 66

R

RAID

RAID 10 overview 22

RAID 5 overview 22

RAID overview 22

recommended configurations for the IBM iSeries 72

recoverability, configuring for 206, 220

registry, setting the TimeOutValue for Windows
2000 206

registry, setting the TimeOutValue for Windows Server
2003 220

remote fibre-channel boot support for a Windows 2000
host system, Installing 207

remote fibre-channel boot support for a Windows Server
2003 host system, Installing 221

remote mirror and copy

restrictions for open-systems hosts 10

removing persistent reserves, Tru64 UNIX 47

requirement, multipath 69

requirements

for fibre-channel attachment

Hewlett-Packard 51, 65

HP 37

Intel 119, 139

iSeries 69

Linux 119

PRIMEPOWER 31

pSeries 87

RS/6000 87

Sun 180

VMware ESX Server 139

Windows 2000 197

Windows Server 2003 211

for iSCSI router 157

SGI attachment 165

restoring data on storage unit 92

RS/6000 and pSeries AIX host system

installing the host attachment package 85

- RS/6000 and pSeries host system
 - verifying the configuration 87
- RS/6000 host system
 - attaching the storage unit 85
 - attachment considerations
 - fibre-channel 87
 - attachment requirements
 - fibre-channel 87
- RS/6000 host, locating the WWPN 227
- RS/6000 pSeries hosts without HACMP, attaching a storage unit to multiple RISC 91

S

- S/390
 - attaching to a FICON channel 110
 - ESCON host cables 100
 - host systems 15
- S/390 and zSeries
 - cable lengths 101
 - controller images and controls 100
 - data transfer 101
 - directors and channel extenders 101
 - host adapters for ESCON 100
 - host systems 15
 - logical paths 101
 - operating system 15
 - port identification 101
 - TSO commands, port identification for 101
- S/390 host system, attaching the storage unit 99
- SAN changes, iSeries 88
- SAN File System host system, attaching the storage unit 161
- SAN File System metadata node, Configuring for attachment 161
- SAN File System metadata server with Fibre-channel adapters, Attaching to 161
- SAN File System, Attaching to an IBM 161
- SAN Volume Controller host system, attaching the storage unit 163
- saving data on the storage unit 91
- SBUS fibre-channel adapter driver, installing the AMCC 182
- SDD (Subsystem Device Driver), installing 194
- SDD considerations for NAS Gateway 500, Multipath I/O and 81
- secondary path, switching I/O 172
- server
 - restoring data on the storage unit 92
 - saving data on the storage unit 91
- setting the parameters for a Sun host system 193
- setting the TimeOutValue registry for Windows 2000 206
- setting the TimeOutValue registry for Windows Server 2003 220
- setting up ESCON and FICON links 108
- setting up the storage unit
 - OpenVMS 42, 55
- settings 187
 - parameters
 - settings for QLogic QLA2342 190

- settings (*continued*)
 - parameters for AMCC FCE-1473 187
 - parameters for AMCC FCE-6460 187
 - parameters for AMCC FCX-6562 187
 - parameters for AMCC FCX2-6562 187
 - parameters for Emulex LP9002L 34
 - parameters for QLogic QLA2310F 190
 - parameters for QLogic QLA2340 190
 - recommended configuration file parameters for Emulex LP10000 185
 - recommended configuration file parameters for Emulex LP10000DC 185
 - recommended configuration file parameters for Emulex LP9002DC 185
 - recommended configuration file parameters for Emulex LP9002L 185
 - recommended configuration file parameters for Emulex LP9002S 185
 - settings
 - parameters for QLogic QLA2342 190
- SGL
 - confirming storage connectivity 169
 - confirming switch connectivity 167
 - fibre-channel arbitrated loop 170, 172, 175
 - locating the WWPN 229
 - switched fabric 169, 171, 173
- SGL host system
 - attaching the storage unit 165
 - attachment requirements, fibre-channel 165
 - fibre-channel attachment considerations 165
 - PCI-FC-1P-OPT adapter card 165
 - PCI-FC-1P-OPT-A adapter card 165
 - PCI-FC-1P-OPT-B adapter card 165
 - PCX-FC-2P-OPT-B adapter card 165
- short wave, fibre-channel 19
- software
 - IBM iSeries 71
 - S/390 and zSeries 15
- state-change notifications (RSCNs) on zSeries hosts, Registered 111
- storage
 - configuring SGL 172
 - configuring the storage unit for Linux 125
 - configuring, SGL 172
- storage area networks (SANs), fibre-channel 21
- storage connectivity
 - OpenVMS 43, 59
 - Tru64 UNIX 43
- storage server
 - restoring data 92
 - saving data 91
- storage unit
 - configuring storage on a Linux 125
 - creating file systems on Linux 127
 - host systems supported by 13
 - partition for Linux 125
 - restoring data on the 92
 - saving data 91
 - using file systems on Linux 127
- storage unit configuration
 - verifying for a pSeries host system 87

- storage unit configuration (*continued*)
 - verifying for an RS/6000 87
- storage unit volume, displaying 43, 59
- storage unit volumes for Tru64 UNIX, verifying attachment 45
- storage unit, configuring MC/ServiceGuard 66
- storage, configuring 46, 61
- Subsystem Device Driver (SDD), installing 194
- Subsystem Device Driver, Installing Linux over the SAN 132
- Subsystem Device Driver, Installing Linux over the SAN without an IBM 131
- Subsystem Device Driver, Multipathing with the 83
- Sun host system
 - attachment requirements
 - fibre-channel 180
 - kernel, changing 184
 - setting parameters
 - AMCC adapters 193
 - Emulex adapters 193
 - QLogic adapters 193
- Sun host, locating the WWPNN 230
- Support for fibre-channel boot 130
- support for NAS Gateway 500, Multipathing 81
- support for the AlphaServer console 38, 39
- Supported adapter cards for IBM NAS Gateway 500 hosts 79
- switch connectivity
 - OpenVMS 58
 - Tru64 UNIX 42
- switch, displaying zoning information SGI 168
- switched fabric, SGI 169, 171, 173
- switched-fabric topology 16
 - description 16
 - illustration of 16
- switching I/O between the primary and secondary alternate path, SGI 172
- systems
 - S/390 and zSeries host 15

T

- targets and LUNs 20
- TimeOutValue registry for Windows 2000, setting 206
- TimeOutValue registry for Windows Server 2003, setting 220
- topology
 - arbitrated loop 16
 - fibre-channel 16
 - point-to-point 16
 - switched-fabric 16
- Trademarks 237
- transfer, data for S/390 and zSeries 101
- Tru64 UNIX
 - assigning volume 42
 - limitations 50
 - storage connectivity 43
 - switch connectivity 42
 - verifying attachment 45
 - Version 5.0x 40

- TSO commands, port identification for S/390 and zSeries 101

U

- unit
 - controls, for S/390 and zSeries 100
 - images, for S/390 and zSeries 100
- updating
 - Windows 2000 device driver 203
 - Windows Server 2003 device driver 217
- using a Web browser, Obtaining WWPNNs 79
- using file systems on the storage unit, Linux 127

V

- verifying
 - attachment of the storage unit volumes for Tru64 UNIX 45
 - host configured for storage 158, 206, 220
 - storage unit configuration for an RS/6000 and pSeries host system 87
- Version 5 Release 3, i5/OS 69, 70, 73
- VMware ESX host system
 - attaching the storage unit 139
- VMware ESX Server host system
 - attachment considerations, fibre-channel 140
 - attachment requirements, fibre-channel 139
- volumes for Tru64 UNIX, verifying attachment 45

W

- Web browser, Obtaining WWPNNs using a 79
- Windows 2000 device driver, updating the 203
- Windows 2000 host system
 - attaching the storage unit 197
 - attachment considerations, fibre-channel 198
 - attachment requirements
 - fibre-channel 197
 - configuring for availability and recoverability 206
- Windows 2000 host system, Installing remote fibre-channel boot support for a 207
- Windows 2000 host, locating the WWPNN 230
- Windows 2003 host system attachment considerations, fibre-channel 212
- Windows Server 2003 device driver, updating the 217
- Windows Server 2003 host system
 - attaching the storage unit 211
 - attachment requirements
 - fibre-channel 211
 - configuring for availability and recoverability 220
 - Emulex LP10000 adapter card 211
 - Emulex LP10000DC adapter card 211
 - Emulex LP9002L adapter card 211
 - Emulex LP9402DC adapter card 211
 - Emulex LP9802 adapter card 211
 - Netfinity adapter card 211
 - QLogic QLA2310FL adapter card 211
 - QLogic QLA2340FL adapter card 211
 - QLogic QLA2342FL adapter card 211

- Windows Server 2003 host system, Installing remote fibre-channel boot support for a 221
- Windows Server 2003 host, locating the WWPN 231
- worldwide port name 225
- worldwide port name, Finding the 79
- worldwide port name, locating
 - for a Hewlett-Packard host 226
 - for an HP-UX host 226
 - for an iSeries host 227
 - Linux host 228
 - Novell NetWare host 229
 - pSeries host 227
 - RS/6000 host 227
 - SGI host 229
 - Sun host 230
 - Windows 2000 host 230
 - Windows Server 2003 host 231
- WWPN 225
- WWPNs through the command-line interface, Obtaining 79
- WWPNs using a Web browser, Obtaining 79

Z

- zoning switch information, SGI 168
- zSeries
 - attaching to a FICON channel 110
 - ESCON host cables 100
 - host systems 15
 - See S/390 and zSeries 15
- zSeries host running Linux, Attaching a storage unit to an S/390 or 112
- zSeries host system, attaching the storage unit 99
- zSeries host, Accessing the storage unit using FCP for a 114
- zSeries hosts, Registered state-change notifications (RSCNs) on 111

Readers' comments — we would like to hear from you

IBM TotalStorage DS8000
Host Systems Attachment Guide

Publication No. SC26-7628-03

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Cut or Fold
Along Line

Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department 61C
9032 South Rita Road
Tucson, AZ 85775-4401



Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold
Along Line



Printed in USA

SC26-7628-03

