

IBM System Storage

*Multipath Subsystem Device Driver
User's Guide*



IBM License Agreement for Machine Code

This guide might contain references to machine code, which includes Licensed Internal Code. Licensed Internal Code is licensed to you under the terms of the IBM License Agreement for Machine Code. Carefully read the agreement. By using this product, you agree to abide by the terms of this agreement and applicable copyright laws. See "IBM license agreement for machine code" on page 415.

Note

Before using this information and the product it supports, read the information in "Notices" on page 413.

This edition applies to the following versions of IBM Multipath Subsystem Device Driver and to all subsequent releases and modifications until otherwise indicated in new editions:

- Subsystem Device Driver Version 1 Release 8 Modification 0 Level *x* for HP-UX
- Subsystem Device Driver Version 1 Release 7 Modification 2 Level *x* for AIX
- Subsystem Device Driver Version 1 Release 6 Modification 5 Level *x* for Solaris
- Subsystem Device Driver Version 1 Release 6 Modification 3 Level *x* for Linux
- Subsystem Device Driver Version 1 Release 6 Modification 0 Level *x* for Netware
- | • Subsystem Device Driver Device Specific Module Version 2 Release 4 Modification 5 Level 0 for Windows
- Subsystem Device Driver Path Control Module Version 3 Release 0 Modification *x* Level *x*
- Subsystem Device Driver Path Control Module Version 2 Release 6 Modification 6 Level 0

This edition replaces GC52-1309-07.

© Copyright IBM Corporation 1999, 2015.

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

Contents

Figures ix

Tables xi

About this guide xiii

Who should use this guide	xiii
Summary of changes	xiii
Updated information	xiv
Command syntax conventions	xiv
Highlighting conventions	xiv
Special characters conventions	xiv
Related information	xiv
The ESS library	xv
The DS8000 library	xvi
The DS6000 library	xvii
The DS5000 and DS Storage Manager library	xvii
The DS4000 library	xviii
The SAN Volume Controller library	xviii
The Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication libraries	xix
Ordering IBM publications	xix
How to send your comments	xx

Chapter 1. Overview of the SDD 1

The SDD architecture	2
Enhanced data availability	4
Dynamic I/O load balancing	6
Automatic path-failover protection	6
Concurrent download of licensed machine code for disk storage systems	7
Concurrent download of licensed machine code for virtualization products	7
Concurrent download of licensed machine code for DS3950, DS4000 and DS5000	7
Concurrent download of licensed machine code for IBM BladeCenter S SAS RAID Controller Module (RSSM)	8
Active/Passive dual array controller path-selection algorithm for DS3950, DS4000 and DS5000 products	8
Preferred node path-selection algorithm for the virtualization products	8
Open HyperSwap Replication	9

Chapter 2. Using the SDD on an AIX host system 11

Supported SDD features	11
Verifying the hardware and software requirements	12
Hardware	12
Software	12
Unsupported environments	13
Host system requirements	13
Preparing for the SDD installation	15
Configuring the disk storage system	15

Configuring the virtualization products	15
Installing the AIX fibre-channel device drivers	16
Uninstalling the AIX fibre-channel device drivers	17
Installing and upgrading the AIX SDD host attachment	17
Configuring fibre-channel-attached devices	18
Removing fibre-channel-attached devices	18
Verifying the adapter firmware level	19
Determining if the sddServer for Expert is installed	19
Understanding SDD support for IBM System p with static LPARs configured	20
Determining the installation package	20
Determining the installation type	22
Installing and upgrading the SDD	23
Installing the SDD	23
Upgrading the SDD	24
Verifying the currently installed version of SDD for SDD 1.3.3.11 (or earlier)	35
Verifying the currently installed version of SDD for SDD 1.4.0.0 (or later)	36
Preparing to configure SDD	38
Maximum number of LUNs	38
Controlling I/O flow to SDD devices with the SDD qdepth_enable attribute	42
Controlling reserve policy of SDD devices with the SDD reserve_policy attribute	44
Configuring SDD	44
Unconfiguring SDD	45
Verifying the SDD configuration	45
Dynamically adding paths to SDD vpath devices	46
Dynamically removing or replacing PCI adapters or paths	47
Removing SDD from an AIX host system	49
Special considerations when you uninstall SDD in the NIM environment	50
Removing SDD Host Attachment from an AIX host system	50
SAN boot support	51
Manual exclusion of devices from the SDD configuration	51
Replacing manually excluded devices in the SDD configuration	52
SAN boot installation procedures	52
Understanding SDD support for High Availability Cluster Multiprocessing	53
SDD persistent reserve attributes	56
Preparation for importing volume groups under HACMP	56
HACMP RAID concurrent-mode volume groups and enhanced concurrent-capable volume groups	57
SDD server daemon	67
Verifying if the SDD server has started	67
Starting the SDD server manually	67
Changing to a different port number for the SDD server	67

Stopping the SDD server	68
Replacing the SDD server with a stand-alone version	68
PTFs for APARs on AIX with Fibre Channel and the SDD server	68
Dynamically changing the SDD path-selection policy algorithm	69
datapath set device policy command	70
Dynamically opening an invalid or close_dead path	70
Fibre-channel dynamic device tracking for AIX 5.20 TL1 (and later)	70
Understanding SDD 1.3.2.9 (or later) support for single-path configuration for supported storage devices	71
Understanding the persistent reserve issue when migrating from SDD to non-SDD volume groups after a system reboot	71
Managing secondary-system paging space	72
Listing paging spaces	72
Adding a paging space	72
Removing a paging space	72
Verifying load-balancing and failover protection	73
Displaying the supported storage device SDD vpath device configuration	73
Configuring volume groups for failover protection	74
Losing failover protection	75
Losing a device path	76
Creating a volume group from single-path SDD vpath devices.	76
A side effect of running the disk change method	76
Manually deleting devices and running the configuration manager (cfgmgr)	78
Using LVM with SDD	78
Importing volume groups with SDD	78
Exporting a volume group with SDD	79
Recovering from mixed volume groups	80
Extending an existing SDD volume group	80
Backing up all files belonging to an SDD volume group	80
Restoring all files belonging to an SDD volume group	81
SDD-specific SMIT panels	81
Accessing the Display Data Path Device Configuration SMIT panel	82
Accessing the Display Data Path Device Status SMIT panel	83
Accessing the Display Data Path Device Adapter Status SMIT panel	83
Accessing the Define and Configure All Data Path Devices SMIT panel	84
Accessing the Add Paths to Available Data Path Devices SMIT panel	84
Accessing the Configure a Defined Data Path Device SMIT panel	84
Accessing the Remove a Data Path Device SMIT panel	84
Accessing the Add a Volume Group with Data Path Devices SMIT panel	84
Accessing the Add a Data Path Volume to a Volume Group SMIT panel	85

Accessing the Remove a Physical Volume from a Volume Group SMIT panel	85
Accessing the Backup a Volume Group with Data Path Devices SMIT panel	85
Accessing the Remake a Volume Group with Data Path Devices SMIT panel	86
SDD utility programs	86
addpaths	86
hd2vp and vp2hd	87
dpovgfix	87
lsvpcfg	87
mkvg4vp	88
extendvg4vp	89
excludesddcfg	89
lquerypr	90
sddgetdata	90
Persistent reserve command tool	90
Using supported storage devices directly	92
Using supported storage devices through AIX LVM	93
Migrating a non-SDD volume group to a supported storage device SDD multipath volume group in concurrent mode	93
Detailed instructions for migrating a non-SDD volume group to a supported storage device SDD multipath volume group in concurrent mode	95
Using the trace function	97

Chapter 3. Using SDDPCM on an AIX host system 99

Supported SDDPCM features	101
Unsupported SDDPCM features	102
Verifying the hardware and software requirements	103
Hardware	103
Software	103
Unsupported environments.	103
Host system requirements	104
Preparing for SDDPCM installation	105
Preparing for SDDPCM installation for supported storage devices	106
Installing and upgrading the AIX SDDPCM host attachment	109
Installing and upgrading SDDPCM	111
Installing SDDPCM from CD-ROM	111
Installing SDDPCM from downloaded code	113
Installing SDDPCM with the AIX OS from an AIX NIM SPOT server to the client SAN boot disk or the internal boot disk	114
Updating SDDPCM	115
Migrating SDDPCM	117
Configuring and unconfiguring supported storage MPIO-capable devices	120
Configuring supported storage MPIO-capable devices	120
Verifying the SDDPCM configuration	120
Dynamically adding and removing paths or adapters	121
Unconfiguring supported storage MPIO-capable devices	122
Switching DS4000, DS5000, or DS3950 storage device configurations.	123
Removing SDDPCM from an AIX host system	123

SDDPCM support for HACMP with Enhanced Concurrent Mode volume groups	124
MPIO reserve policies	125
No Reserve reservation policy	125
Exclusive Host Access single-path reservation policy	125
Persistent Reserve Exclusive Host Access reservation policy	126
Persistent Reserve Shared Host Access reservation policy	126
SDDPCM active/passive storage device controller health-check feature	126
SDDPCM ODM attribute settings	127
SDDPCM ODM attribute default settings	127
Changing device reserve policies	127
Changing the path selection algorithm	128
Using the load balancing port algorithm	128
Changing SDDPCM path healthcheck mode	130
Changing SDDPCM path healthcheck time interval	131
Changing the Open HyperSwap quiesce expire time	131
Supported AIX fibre-channel device driver features	132
Fast I/O failure of fibre-channel devices	132
Fibre channel dynamic device tracking	133
Changing SDDPCM controller healthcheck delay_time	133
Changing SDDPCM controller healthcheck interval time	134
Multipath SAN boot support	135
Configuring supported storage system MPIO devices as the SAN boot device	135
Support system dump device with the supported storage system MPIO device	136
Dynamically enabling and disabling paths or adapters	136
Dynamically enabling or disabling a path	136
Dynamically enabling or disabling an adapter	137
Using the SDDPCM trace function	138
SDDPCM server daemon	139
Verifying if the SDDPCM server has started	139
Starting the SDDPCM server manually	139
Stopping the SDDPCM server	140
AE daemon	140
Verifying if the AE server has started	140
Starting the AE server manually	141
Stopping the AE server manually	141
SDDPCM utility programs	141
Persistent reserve command tools	141
Using SDDPCM pcmpath commands	145

Chapter 4. Using the SDD on an HP-UX host system. 175

Verifying the hardware and software requirements	175
Hardware	175
Software	175
Unsupported environments.	176
Understanding how SDD works on an HP-UX host system	176
Preparing for the SDD installation	176
Configuring the disk storage system.	176

Configuring the virtualization products	176
Configuring DS4000 and DS5000 family disk systems	176
Bypassing the boot disk check.	177
Planning for installation.	177
Determining if the SDD 1.3.1.5 (or later) server for Expert is installed.	178
Installing SDD	179
Installing SDD from CD-ROM.	179
Installing SDD from downloaded code	180
Upgrading the SDD	181
Installing an additional package for DS4000 and DS5000 support	183
Configuring the SDD.	183
Maximum number of LUNs	183
Changing an SDD hardware configuration	183
Converting a volume group	184
Dynamic reconfiguration	184
Dynamically changing the SDD path-selection policy algorithm	184
Preferred node path-selection algorithm for the virtualization products	185
SDD datapath query adapter command changes for SDD 1.4.0.0 (or later).	186
SDD datapath query device command changes for SDD 1.4.0.0 (or later).	187
Postinstallation	188
Uninstalling SDD	190
SDD server daemon	191
Verifying if the SDD server has started	191
Starting the SDD server manually	192
Changing to a different port number for the SDD server	192
Stopping the SDD server	192
How to import and export volume groups	192
Exporting volume groups	193
Moving the map file	193
Creating the volume group device directory	193
Creating the group special file.	194
Importing volume groups	194
Exporting or importing MC Service Guard volume groups	194
Using applications with SDD	196
Standard UNIX applications	196
Installing SDD on a NFS file server	201

Chapter 5. Using SDD on a Linux host system 203

Verifying hardware and software requirements	203
Hardware	203
Software	203
Unsupported environments.	204
Preparing for SDD installation.	204
Configuring disk storage systems.	204
Configuring virtualization products	204
Configuring fibre-channel adapters on disk storage systems.	204
Configuring fibre-channel adapters on virtualization products	205
Disabling automatic Linux system updates	205
Installing SDD	205

Installing SDD from CD-ROM	205
Installing SDD from downloaded code	206
Upgrading SDD	206
Verifying the SDD installation	207
Configuring SDD	208
Configuring and verifying an SDD	209
Reconfiguring using SDD userspace commands	213
Configuring SDD at system startup	214
Maintaining SDD vpath device configuration persistence	215
Dynamically changing the SDD path-selection policy algorithm	216
Using dynamic reconfiguration	217
Uninstalling SDD	217
Mounting SDD devices with automount	218
Setting up automount	218
Configuring automount	218
Sample mount output	220
Using Linux Logical Volume Manager with SDD	221
Using LVM on SUSE	221
Using LVM on Red Hat	223
LVM 2 general configuration information	223
Booting Linux over the SAN with SDD	224
SAN Boot instructions for RHEL 3 with IBM SDD (x86)	225
SAN Boot instructions for RHEL 3 with IBM SDD (ppc)	232
SAN Boot Instructions for SLES 8 with IBM SDD (x86)	240
SAN Boot Instructions for SLES 9 with IBM SDD (x86)	249
SAN Boot instructions for SLES 9 with IBM (ppc)	255
SAN Boot Instructions for SLES 9 with IBM SDD (x86) and LVM 2	262
SAN boot instructions for RHEL 4 with IBM SDD (x86)	270
SAN Boot instructions for RHEL 4 with IBM SDD (ppc)	276
SAN boot instructions for RHEL 4 with IBM SDD (x86) and LVM 2	282
Using lilo with SDD (remote boot) on x86	289
SDD server daemon	292
Verifying if the SDD server has started	292
Starting the SDD server manually	292
Changing to a different port number for the SDD server	293
Stopping the SDD server	293
Collecting trace information	293
Understanding SDD support for single-path configuration	293
Partitioning SDD vpath devices	294
Using standard UNIX applications	294
Managing common issues	295

Chapter 6. Using the SDD on a NetWare host system 297

Verifying the hardware and software requirements	297
Hardware requirements	297
Software requirements	297
Supported environments	298

Unsupported environments.	298
Disk storage system requirements	298
SCSI requirements.	298
Fibre-channel requirements.	299
Preparing for SDD installation.	299
Configuring the disk storage system.	299
Configuring fibre-channel adapters	299
Configuring SCSI adapters	300
Using a NetWare Compaq Server.	300
Installing SDD	301
Installing SDD from CD-ROM.	301
Installing SDD from downloaded code	301
Configuring SDD	301
Maximum number of LUNs	302
Displaying the current version of the SDD.	302
Features	302
Automatic path detection, failover and selection	302
Manual operations using the datapath commands	303
Understanding SDD error recovery algorithms	304
Dynamic load balancing.	304
Disk storage system logical unit detection	304
Error reporting and logging	304
SDD in NetWare-layered architecture	305
Removing the SDD	305
Cluster setup for Novell NetWare 5.1	306
Cluster setup for Novell NetWare 6.0	306
Examples of commands output on the Console Window	306

Chapter 7. Using the SDD on a Solaris host system 311

Verifying the hardware and software requirements	311
Hardware	311
Software	311
Supported environments	312
Unsupported environments.	312
Understanding how SDD works on a Solaris host system	312
Preparing for SDD installation.	312
Configuring disk storage systems.	312
Configuring virtualization products	312
Determining if the SDD server for Expert is installed	313
Planning for installation.	313
Installing SDD	315
Installing SDD from CD-ROM.	315
Installing SDD from downloaded code	317
Postinstallation	317
Verifying the SDD installation	320
Upgrading the SDD	320
Configuring the SDD.	321
Changing an SDD hardware configuration	321
Solaris 10 ZFS support	323
Solaris 10 Zone support	323
Dynamically changing the SDD path-selection policy algorithm	324
Excluding LUNs from being configured by SDD	325
Uninstalling the SDD.	326
Understanding SDD support for single-path configuration for disk storage system	327

SDD server daemon	327
Verifying if the SDD server has started	327
Starting the SDD server manually	328
Changing to a different port number for the SDD server	328
Changing the retry count value when probing SDD server inquiries	328
Stopping the SDD server	328
Using applications with SDD	329
Standard UNIX applications	329
Installing the SDD on a NFS file server	329
Veritas Volume Manager	330
Oracle	331
Solaris Volume Manager (formerly Solstice DiskSuite)	335
Chapter 8. Using SDDDSM on a Windows Server 2008 or Windows Server 2012 host system	339
Verifying the hardware and software requirements	339
Unsupported environments	340
Host system requirements	340
Preparing for SDDDSM installation	340
Configuring the supported storage device	341
Configuring fibre-channel adapters	341
Installing SDDDSM	341
Installing for the first time	341
Upgrading SDDDSM	343
Displaying the current version of SDDDSM	343
Configuring SDDDSM	344
Maximum number of LUNs	344
Verifying the configuration	344
Activating additional paths	345
Verifying that additional paths are installed correctly	345
Removing additional paths	347
Verifying that additional paths are removed correctly	347
Uninstalling SDDDSM	350
SAN boot support	351
Remote boot support for 32-bit Windows Server 2008 or Windows Server 2012 using a QLogic HBA	351
Booting from a SAN device with Windows Server 2008 or Windows Server 2012 and the SDDDSM using an Emulex HBA	352
Support for Windows Server 2008 or Windows Server 2012 clustering	353
Configuring a Windows Server 2008 or Windows Server 2012 cluster with SDDDSM installed	353
Removing SDDDSM in a two-node cluster environment	354
About I/O timeout errors	354
SDDDSM datapath command support	354
SDDDSM server daemon	355
Verifying if the SDDDSM server has started	355
Starting the SDDDSM server manually	355
Changing to a different port number for the SDDDSM server	356

Stopping the SDDDSM server	356
SDDDSM 2.4.5.0 features	356
Additional load balance policies in SDDDSM 2.4.5.0	356
iSCSI support	358
iSCSI in Hyper-V support	359
FC in Hyper-V support	360
Path Performance Plug-in	361
Enhancement of failover group limit per adapter	363

Chapter 9. Using the SDDDSM or the SDDPCM daemon 365

SDD server daemon	365
Understanding how the SDD server daemon works	365
SDDPCM server daemon	367
sddsrv.conf file format	367
pcmsrv.conf file format	368
Enabling or disabling the sddsrv or pcmsrv TCP/IP port	369
Changing the sddsrv or pcmsrv TCP/IP port number	369
Disabling the probing function in sddsrv	369
Changing the probing interval in sddsrv	370

Chapter 10. Using the datapath commands 371

datapath clear device count	372
datapath disable ports	373
datapath enable ports	374
datapath open device path	374
datapath query adapter	376
datapath query adaptstats	377
datapath query device	378
datapath query devstats	381
datapath query essmap	383
datapath query portmap	384
datapath query version	385
datapath query wwpn	386
datapath remove adapter	386
datapath remove device path	387
datapath set adapter	389
datapath set device policy	389
datapath set device path	390
datapath set performance	391
datapath set qdepth	391

Appendix A. SDD, SDDPCM, and SDDDSM data collection for problem analysis. 393

Enhanced trace capability for the SDD and SDDDSM	393
Using sddgetdata to collect information for problem determination	393
Enhanced trace capability for SDDPCM	394
Using sddpcmgetdata to collect information for problem determination	395

Appendix B. System log messages	397
AIX error and information messages	397
SDDPCM error log messages	398
Windows log messages	400
Appendix C. SDDDSM tools	405
Notices	413
Trademarks	415

IBM license agreement for machine code	415
Glossary	419
Index	449

Figures

1. Multipath connections between a host system and the disk storage in a disk storage system 5
2. Multipath connections between a host system and the disk storage with the SAN Volume Controller 6
3. SDDPCM in the protocol stack 100
4. Workload imbalance when one link receives twice the load of the other links 129
5. Workload imbalance when one link is more heavily loaded than another link 129
6. Workload imbalance when one host sharing workload across two paths loses one path 130
7. Example showing SAN Volume Controller devices to the host and path access to the SAN Volume Controller devices in a successful SDDDSM installation on a Windows Server host system 342

Tables

1.	SDD platforms on supported storage devices	1	12.	PTFs for APARs on AIX with fibre-channel support and the SDD server daemon running	68
2.	SDD in the protocol stack	3	13.	SDD-specific SMIT panels and how to proceed	81
3.	Package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0 (or later)	12	14.	Commands	145
4.	SDD 1.4.0.0 (or later) installation packages for different AIX OS levels and the supported AIX kernel mode, application mode, and interface	20	15.	SDD installation scenarios	177
5.	Major files included in the SDD installation package	21	16.	Patches necessary for proper operation of SDD on HP-UX	178
6.	List of previously installed installation packages that are supported with the installation upgrade	24	17.	SDD components installed for HP-UX host systems	188
7.	Maximum LUNs allowed for different AIX OS levels	39	18.	System files updated for HP-UX host systems	188
8.	Recommended maximum paths supported for different number of LUNs on AIX 5.2 or later	39	19.	SDD commands and their descriptions for HP-UX host systems	189
9.	Recommended SDD installation packages and supported HACMP modes for SDD versions earlier than SDD 1.4.0.0	54	20.	SDD components for a Linux host system	207
10.	Software support for HACMP 4.5 on AIX 4.3.3 (32-bit only), 5.1.0 (32-bit and 64-bit), 5.2.0 (32-bit and 64-bit)	55	21.	Summary of SDD commands for a Linux host system	208
11.	Software support for HACMP 4.5 on AIX 5.1.0 (32-bit and 64-bit kernel)	55	22.	SDD installation scenarios	314
			23.	Operating systems and SDD package file names	314
			24.	SDD components installed for Solaris host systems	317
			25.	System files updated for Solaris host systems	318
			26.	SDD commands and their descriptions for Solaris host systems	319
			27.	Commands	371

About this guide

The IBM® System Storage® Multipath Subsystem Device Driver (SDD) supports a storage-redundant configuration environment for a host system that is attached to storage devices. It provides enhanced data availability, dynamic input/output (I/O) load balancing across multiple paths, and automatic path-failover protection for the following host systems:

- IBM AIX®
- HP-UX
- Supported Linux distributions, levels, and architectures
- Novell Netware
- Sun Solaris
- Microsoft Windows Server 2008
- Microsoft Windows Server 2012

The IBM System Storage Multipath Subsystem Device Driver Path Control Module (SDDPCM) provides AIX Multipath I/O (MPIO) support. It is a loadable module. During the configuration of supported storage devices, SDDPCM is loaded and becomes part of the AIX MPIO Fibre Channel protocol device driver. The AIX MPIO-capable device driver with the SDDPCM module provides the same functions that SDD provides.

The IBM System Storage Multipath Subsystem Device Driver Device Specific Module (SDDDSM) provides multipath I/O support for Windows platforms based on the MPIO technology provided by Microsoft.

Who should use this guide

This guide is intended for users who will install and configure IBM System Storage Multipath Subsystem Device Driver (SDD), IBM System Storage Multipath Subsystem Device Driver Path Control Module (SDDPCM), or IBM System Storage Multipath Subsystem Device Driver Device Specific Module (SDDDSM). The installer of SDD, SDDPCM, or SDDDSM should be:

- Familiar with the operating systems on which the driver will be installed
- Familiar with the storage systems on which the driver will be operated
- Familiar with the applications used with the driver
- Aware of the procedures in this document
- Aware of related installation and service publications

Summary of changes

This guide contains information that was published in the *IBM System Storage Multipath Subsystem Device Driver User's Guide* and technical updates in that information. All changes to this guide are marked with a vertical bar (|) in the left margin.

Note: For the more recent updates that are not included in this guide, go to the SDD website at: www-947.ibm.com/support/entry/portal/product/system_storage/storage_software/other_software_products/system_storage_multipath_subsystem_device_driver?productContext=-1564959213

Updated information

This edition of the guide includes information about the following updates:

- New features added with SDDDSM 2.4.5.0. For more information, see “SDDDSM 2.4.5.0 features” on page 356.

Command syntax conventions

This section describes the notational conventions that this book uses.

Highlighting conventions

The following typefaces are used to show emphasis:

boldface

Text in **boldface** represents menu items and 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 commands that you type, samples of command output, examples of program code or messages from the system, and configuration state of the paths or volumes (such as Dead, Active, Open, Closed, Online, Offline, Invalid, Available, Defined).

Special characters conventions

The following special character conventions are used in this book:

*** asterisks**

Asterisks (*) are used as wildcard symbols to search for the beginning or remaining characters of the installation package name.

For example, the asterisks in the beginning and at the end of Sdd characters in the `ls1pp -l *Sdd*` command are used as wildcard symbols to search for the characters `ibm...` and `...rte`.

... ellipsis

Ellipsis indicates that more commands are in the next command line.

`< >` Indicate optional parameters.

Related information

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

- The publications for the IBM System Storage DS8000[®] library
- The publications for the IBM System Storage DS6000[™] library
- The publications for the IBM System Storage DS5000 and DS Storage Manager library
- The publications for the IBM System Storage DS4000[®] library
- The publications for the IBM System Storage SAN Volume Controller library
- The publications for the IBM Tivoli[®] Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication libraries
- The publications for the IBM TotalStorage Enterprise Storage Server[®] (ESS) library
- Other IBM publications that relate to the ESS
- Non-IBM publications that relate to the ESS

Note: IBM does not support ESS. For more information, contact the IBM Support team.

See “Ordering IBM publications” on page xix for information about how to order publications. See “How to send your comments” on page xx for information about how to send comments about the publications.

The ESS library

Note: IBM does not support ESS. For more information, contact the IBM Support team.

The following customer publications make up the ESS library. Unless otherwise noted, these publications are available in Adobe portable document format (PDF) on a compact disc (CD) that comes with the ESS. If you need additional copies of this CD, the order number is SK2T-8803.

See “IBM Publications Center” on page xix for information about ordering these and other IBM publications.

Title	Description	Order number
IBM TotalStorage Enterprise Storage Server Copy Services Command-Line Interface Reference	This guide describes the commands that you can use from the ESS Copy Services command-line interface (CLI) for managing your ESS 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 an ESS Copy Services server application. You can use the CLI commands to indirectly control Peer-to-Peer Remote Copy (PPRC) and IBM FlashCopy® configuration tasks within an ESS Copy Services server group.	SC26-7494 (See Note.)
IBM TotalStorage Enterprise Storage Server Configuration Planner for Open-Systems Hosts	This guide provides guidelines and work sheets for planning the logical configuration of an ESS that attaches to open-systems hosts.	SC26-7477 (See Note.)
IBM TotalStorage Enterprise Storage Server Configuration Planner for S/390® and IBM eServer™ zSeries Hosts	This guide provides guidelines and work sheets for planning the logical configuration of an ESS that attaches to either the IBM S/390 or IBM eServer zSeries host system.	SC26-7476 (See Note.)
IBM TotalStorage Enterprise Storage Server Host Systems Attachment Guide	This guide provides guidelines for attaching the ESS to your host system and for migrating to fibre-channel attachment from either a Small Computer System Interface (SCSI) or from the IBM SAN Data Gateway.	SC26-7446 (See Note.)
IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide	This guide introduces the ESS product and lists the features you can order. It also provides guidelines for planning the installation and configuration of the ESS.	GC26-7444

Title	Description	Order number
IBM TotalStorage Storage Solutions Safety Notices	This publication provides translations of the danger notices and caution notices that IBM uses in ESS publications.	GC26-7229
IBM TotalStorage Enterprise Storage Server SCSI Command Reference	This publication describes the functions of the ESS. It provides reference information, such as channel commands, sense bytes, and error recovery procedures for UNIX, IBM Application System/400 (IBM AS/400), and IBM eServer iSeries 400 hosts.	SC26-7297
IBM TotalStorage Enterprise Storage Server Subsystem Device Driver User's Guide	This publication describes how to use the IBM TotalStorage ESS Subsystem Device Driver (SDD) on open-systems hosts to enhance performance and availability on the ESS. SDD creates redundant paths for shared LUNs. SDD permits applications to run without interruption when path errors occur. It balances the workload across paths, and it transparently integrates with applications.	SC26-7637
IBM TotalStorage Enterprise Storage Server User's Guide	This guide provides instructions for setting up and operating the ESS and for analyzing problems.	SC26-7445 (See Note.)
IBM TotalStorage Enterprise Storage Server Web Interface User's Guide	This guide provides instructions for using the two ESS Web interfaces: ESS Specialist and ESS Copy Services.	SC26-7448 (See Note.)
IBM TotalStorage Common Information Model Agent for the Enterprise Storage Server Installation and Configuration Guide	This guide introduces the common interface model (CIM) concept and provides instructions for installing and configuring the CIM agent. The CIM agent acts as an open-system standards interpreter, providing a way for other CIM-compliant storage resource management applications (IBM and non-IBM) to interoperate with each other.	GC35-0485
IBM TotalStorage Enterprise Storage Server Application Programming Interface Reference	This reference provides information about the ESS application programming interface (API).	GC35-0489
Note: No hardcopy book is produced for this publication.		

The DS8000 library

The following publications make up the IBM System Storage DS8000 library. These publications are available from www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

Title	Order number
<i>IBM System Storage DS8000 Introduction and Planning Guide</i>	GC35-0515
<i>IBM System Storage DS[®] Command-Line Interface User's Guide for the DS6000 series and DS8000 series</i>	GC53-1127

Title	Order number
<i>IBM System Storage DS8000 Host Systems Attachment Guide</i>	SC26-7917
<i>IBM System Storage DS Application Programming Interface 5.4.1 and 5.4.2 Installation and Reference</i>	GC35-0516
<i>IBM CIM Agent for DS Open Application Programming Interface 5.5</i>	GC26-7929

The DS6000 library

The following publications make up the IBM System Storage DS6000 library. These publications are available from www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

Title	Order number
<i>IBM System Storage DS6000 Installation, Troubleshooting, and Recovery Guide</i>	GC26-7925
<i>IBM System Storage DS6000 Introduction and Planning Guide</i>	GC26-7924
<i>IBM System Storage DS6000 Host System Attachment Guide</i>	GC26-7923
<i>IBM System Storage DS6000 Messages Reference</i>	GC26-7920
<i>IBM System Storage DS Command-Line Interface User's Guide for the DS6000 series and DS8000 series</i>	GC53-1127
<i>IBM System Storage DS6000 Quick Start Guide</i>	GC26-7921

The DS5000 and DS Storage Manager library

The following publications make up the IBM System Storage DS5000 and DS Storage Manager library. These publications are available from www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

Title	Order number
<i>IBM System Storage DS Storage Manager Version 10.x Installation and Host Support Guide</i>	GC53-1135
<i>IBM System Storage DS Storage Manager Version 10.x Copy Services Guide</i>	GC53-1136
<i>IBM System Storage DS4000/DS5000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview</i>	GC53-1137
<i>IBM System Storage DS5100 and DS5300 Installation, User's and Maintenance Guide</i>	GC53-1140
<i>IBM System Storage EXP5000 Storage Expansion Enclosure Installation, User's, and Maintenance Guide</i>	GC53-1141
<i>IBM System Storage DS Storage Manager Command-Line Programming Guide</i>	GC52-1275
<i>IBM System Storage DS5000 Quick Start Guide: Quick Reference for the DS5100, DS5300 and EXP5000</i>	GC53-1134
<i>IBM TotalStorage DS4300 Fibre Channel Storage Subsystem Installation, User's, and Maintenance Guide</i>	GC26-7722

The DS4000 library

The following publications make up the IBM System Storage DS4000 library. These publications are available from www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

Title	Order number
<i>IBM System Storage DS4000 Storage Manager Concepts Guide</i>	GC26-7734
<i>IBM System Storage DS4000 Storage Manager 10 Installation and Host Support Guide</i>	GA76-0422
<i>IBM System Storage DS4000 Storage Manager Copy Services Guide</i>	GC27-2172
<i>IBM System Storage DS4000 Storage Manager Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview</i>	GC26-7907
<i>IBM System Storage DS4000 Hard Drive and Storage Expansion Enclosure Installation and Migration Guide</i>	GC26-7849
<i>IBM System Storage DS3000/DS4000 Command-Line Programming Guide</i>	GC52-1275
<i>IBM System Storage DS4000 EXP420 Storage Expansion Unit Installation, User's and Maintenance Guide</i>	GC27-2050
<i>IBM System Storage DS4000 EXP810 Storage Expansion Enclosure Installation, User's and Maintenance Guide</i>	GC26-7798
<i>IBM TotalStorage DS4000 EXP700 and EXP710 Storage Expansion Enclosures Installation, User's, and Maintenance Guide</i>	GC26-7735
<i>IBM System Storage DS4200/DS4700 Quick Start Guide</i>	GC27-2147
<i>IBM System Storage DS4700 Installation, User's and Maintenance Guide</i>	GC26-7843
<i>IBM System Storage DS4800 Quick Start Guide</i>	GC27-2148
<i>IBM System Storage DS4800 Installation, User's and Maintenance Guide</i>	GC26-7845
<i>IBM System Storage DS4800 Controller Cache Upgrade Kit Instructions</i>	GC26-7774

The SAN Volume Controller library

The following publications make up the SAN Volume Controller library. Unless otherwise noted, these publications are available in Adobe PDF from the following website:

www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss

Title	Order number
<i>IBM System Storage SAN Volume Controller Planning Guide</i>	GA32-0551
<i>IBM System Storage SAN Volume Controller Model 2145-8A4 Hardware Installation Guide</i>	GC27-2219

Title	Order number
<i>IBM System Storage SAN Volume Controller Model 2145-CF8 Hardware Installation Guide</i>	GC52-1356
<i>IBM System Storage SAN Volume Controller Software Installation and Configuration Guide</i>	SC23-6628
<i>IBM System Storage SAN Volume Controller Command-Line Interface User's Guide</i>	SC26-7903
<i>IBM System Storage SAN Volume Controller CIM Agent Developer's Guide</i>	SC23-6665
<i>IBM System Storage SAN Volume Controller Host Attachment Guide</i>	SC26-7905
<i>IBM System Storage SAN Volume Controller Troubleshooting Guide</i>	GC27-2227
<i>IBM System Storage SAN Volume Controller Hardware Maintenance Guide</i>	GC27-2226

The Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication libraries

The following publications make up the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication libraries. These publications are available from the following website:

www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss

Title	Order number
<i>IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide</i>	SC27-2337
<i>IBM Tivoli Storage Productivity Center User's Guide</i>	SC27-2338
<i>IBM Tivoli Storage Productivity Center Messages</i>	SC27-2340
<i>IBM Tivoli Storage Productivity Center Command-Line Interface Reference</i>	SC27-2339
<i>IBM Tivoli Storage Productivity Center Problem Determination Guide</i>	GC27-2342
<i>IBM Tivoli Storage Productivity Center Workflow User's Guide</i>	SC27-2341

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 at no 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 website:

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this book, you can submit them in one of the following ways:

- E-mail: 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 or fax your comments to:

International Business Machines Corporation
RCF Processing Department
Department 61C
9032 South Rita Road
Tucson, Arizona 85775-4401
U.S.A.

Chapter 1. Overview of the SDD

The IBM System Storage Multipath Subsystem Device Driver (SDD) supports a storage-redundant configuration environment for a host system that is attached to the following devices:

- IBM TotalStorage Enterprise Storage Server (ESS)
- IBM System Storage DS8000
- IBM System Storage DS6000
- IBM System Storage SAN Volume Controller
- IBM System Storage DS5000
- IBM System Storage DS4000
- IBM BladeCenter S SAS RAID Controller Module (RSSM)
- IBM System Storage DS3950

Important: ESS and RSSM are not supported by IBM. For any information about ESS or RSSM, contact the IBM Support team.

This guide uses the following terminology:

- The phrase *supported storage devices* refers to the following types of devices:
 - DS3950, DS4100 (AIX only), DS4200, DS4300, DS4500, DS4700, DS4800, DS5020, DS5100, DS5300, DS6000, and DS8000
 - SAN Volume Controller
 - ESS
 - RSSM
- The phrase *disk storage system* refers to ESS, DS8000, or DS6000 devices.
- The phrase *virtualization product* refers to the SAN Volume Controller. Table 1 indicates the products that different SDD platforms support.
- The phrase *DS4000* refers to DS4100 (AIX only), DS4200, DS4300, DS4500, DS4700, and DS4800 devices.
- The phrase *DS5000* refers to DS5100 and DS5300 devices.
- The phrase *RSSM* refers to IBM BladeCenter S SAS RAID Controller Module devices.
- The phrase *Open HyperSwap[®]* refers to Open HyperSwap replication.
- The phrase *Open HyperSwap device* refers to pair of volumes that are managed in a Tivoli Productivity Center for Replication copy set.
- The phrase *Open HyperSwap session* refers to a collection of Tivoli Productivity Center for Replication managed copy sets.

Table 1. SDD platforms on supported storage devices

Supported storage device								
Platform	ESS	DS8000	DS6000	DS5000	DS4000	DS3950	RSSM	SAN Volume Controller
AIX SDD	✓	✓	✓					✓
AIX SDDPCM	✓	✓	✓	✓	✓	✓	✓	✓

Table 1. SDD platforms on supported storage devices (continued)

Platform	Supported storage device							SAN Volume Controller
	ESS	DS8000	DS6000	DS5000	DS4000	DS3950	RSSM	
HP	✓	✓	✓	✓	✓	✓		✓
Linux	✓	✓	✓					✓
Novell	✓	✓	✓					
SUN	✓	✓	✓					✓
Windows 2008 and Windows 2012 SDDDSM	✓	✓	✓				✓	✓

The SDD supports a storage-redundant configuration environment for a host system that is attached to storage devices. It provides enhanced data availability, dynamic input/output (I/O) load-balancing across multiple paths, and automatic path failover protection.

This guide provides step-by-step procedures on how to install, configure, and use SDD features on the following host systems:

- IBM AIX (SDD and SDDPCM)
- HP-UX
- Supported Linux distributions, levels, and architectures. For up to date information about specific kernel levels supported in this release, see the Readme file on the CD-ROM or visit the SDD website:
www-947.ibm.com/support/entry/portal/product/system_storage/storage_software/other_software_products/system_storage_multipath_subsystem_device_driver?productContext=-1564959213
- Novell Netware (disk storage systems only)
- Sun Solaris
- Microsoft Windows 2008 or Windows Server 2012 SDDDSM

The SDD architecture

The SDD is a software solution to support the multipath configuration environments in supported storage devices. It resides in a host system with the native disk device driver and provides the following functions:

- Enhanced data availability
- Dynamic input/output (I/O) load-balancing across multiple paths
- Automatic path failover protection
- Concurrent download of licensed machine code

Table 2 on page 3 shows the position of the SDD in the protocol stack. I/O operations that are sent to the SDD proceed to the host disk driver after path selection. When an active path experiences a failure (such as a cable or controller failure), the SDD dynamically switches to another path.

Table 2. SDD in the protocol stack

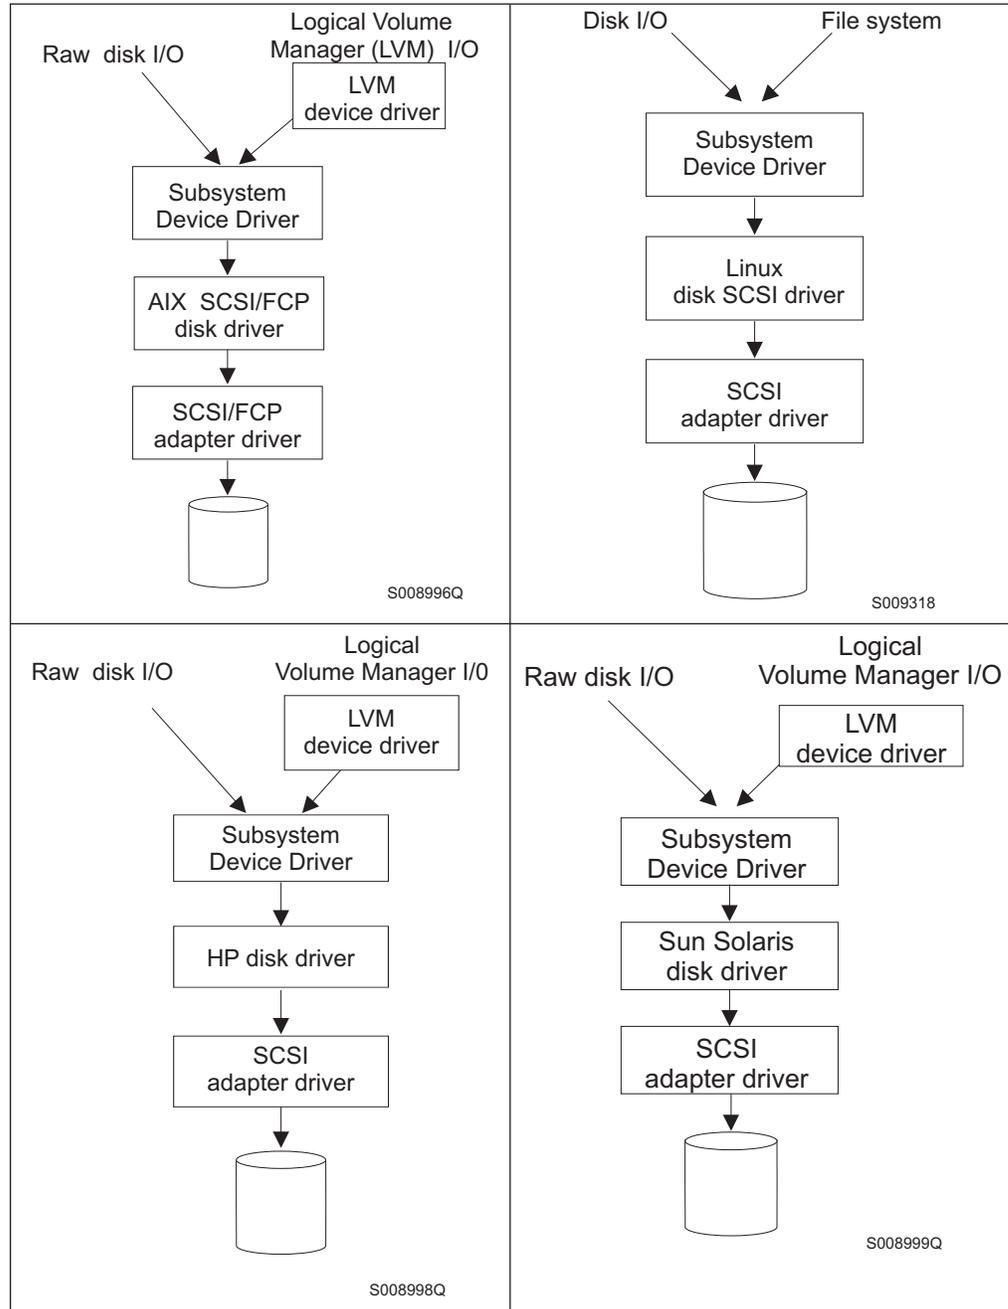
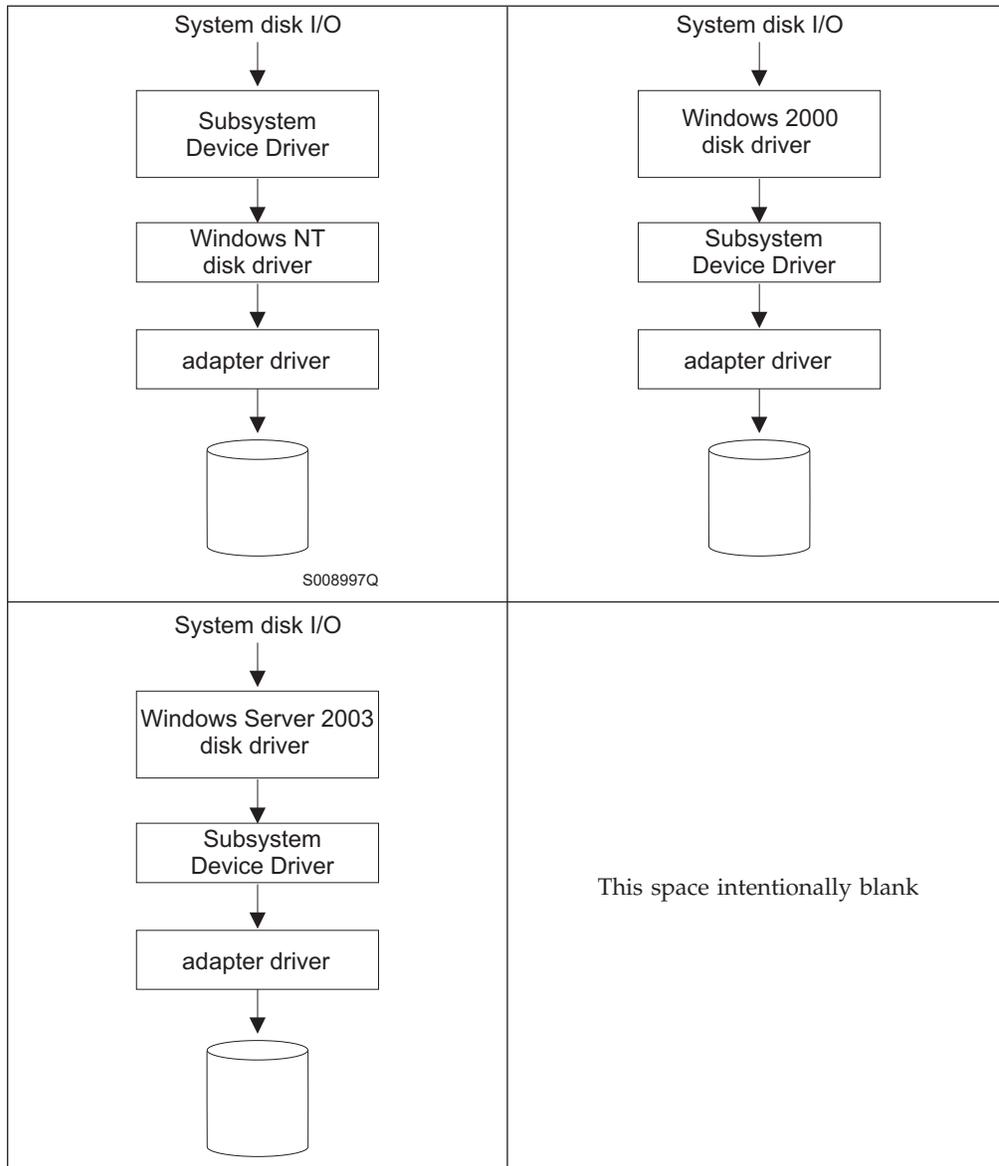


Table 2. SDD in the protocol stack (continued)



Each SDD vpath device represents a unique physical device on the storage server. Each physical device is presented to the operating system as an operating system disk device. There can be up to 32 operating system disk devices that represent up to 32 different paths to the same physical device.

The SDD vpath devices behave almost like native operating system disk devices. You can use most disk device operations of operating systems on the SDD vpath devices, including commands such as **open**, **close**, **dd**, or **fsck**.

Enhanced data availability

Figure 1 on page 5 shows a host system that is attached through small computer system interface (SCSI) or fibre-channel adapters to a disk storage system that has internal component redundancy and multipath configuration. The SDD, residing in the host system, uses this multipath configuration to enhance data availability. That is, when there is a path failure, the SDD reroutes I/O operations from the failing

path to an alternate operational path. This capability prevents a single failing bus adapter on the host system, SCSI or fibre-channel cable, or host-interface adapter on the disk storage system from disrupting data access.

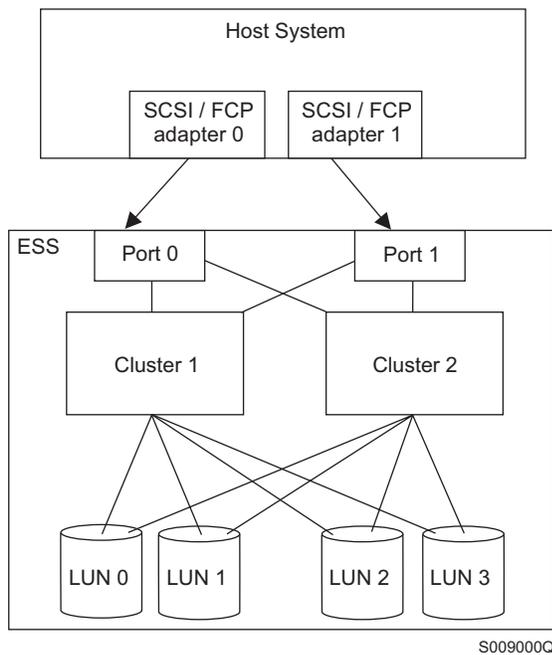


Figure 1. Multipath connections between a host system and the disk storage in a disk storage system

Figure 2 on page 6 shows a host system that is attached through fibre-channel adapters to a SAN Volume Controller that has internal components for redundancy and multipath configuration. The SDD, residing in the host system, uses this multipath configuration to enhance data availability. That is, when there is a path failure, the SDD reroutes I/O operations from the failing path to an alternate operational path. This capability prevents a single failing bus adapter on the host system, fibre-channel cable, or host-interface adapter on the SAN Volume Controller from disrupting data access.

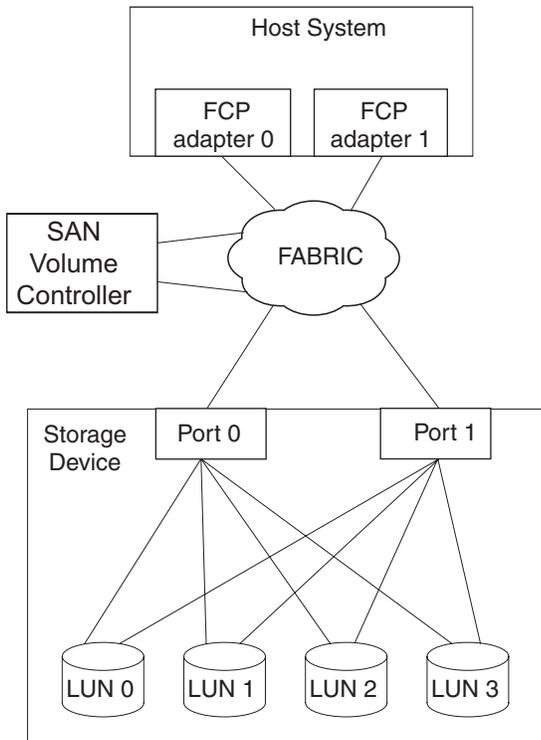


Figure 2. Multipath connections between a host system and the disk storage with the SAN Volume Controller

Note: SAN Volume Controller does not support parallel SCSI attachment.

Dynamic I/O load balancing

By distributing the I/O workload over multiple active paths, the SDD provides dynamic load balancing and eliminates dataflow bottlenecks. In the event of failure in one data path, the SDD automatically switches the affected I/O operations to another active data path, ensuring path-failover protection.

Automatic path-failover protection

The SDD failover protection feature minimizes any disruptions in I/O operations and recovers I/O operations from a failing data path. The SDD provides path-failover protection using the following process:

- Detects a path failure.
- Notifies the host system of the path failure.
- Selects and uses an alternate data path.

The SDD dynamically selects an alternate I/O path when it detects a software or hardware problem.

Some operating system drivers report each detected error in the system error log. With the SDD automatic path-failover feature, some reported errors are actually recovered from an alternative path.

Concurrent download of licensed machine code for disk storage systems

With the SDD multipath mode (configured with at least two paths per multipath device), you can concurrently download and install the licensed machine code (LMC) while applications continue to run. For certain disk storage system LMC, the disk storage system I/O bay or tower are quiesced and resumed. Its adapters might not respond for the duration of the service action, which could be 30 minutes or more.

Note: The SDD does not support single-path mode during the concurrent download of LMC. Also, the SDD does not support single-path mode during any disk storage system concurrent maintenance that impacts the path attachment, such as a disk storage system host-bay-adapter replacement.

Attention: Do not shut down the host or reconfigure the SDD during the concurrent download of licensed machine code or you might lose your initial SDD configuration.

For information about performing the concurrent download of LMC for ESS, see the microcode installation instructions for your specific type and model.

Concurrent download of licensed machine code for virtualization products

With the SDD multipath mode (configured with at least two paths per multipath device), you can concurrently download and install the licensed machine code while your applications continue to run.

At least one path must be configured through each node of a virtualization product group. That is, if only two paths exist, they must go to separate nodes for each I/O group. However, at least two paths to each node are recommended.

During the code upgrade, each node of an I/O group is upgraded sequentially. The node that is being upgraded is temporarily unavailable, and all I/O operations to that node fail. However, failed I/O operations are directed to the other node of the I/O group, and applications do not see any I/O failures.

Attention: Do not shut down the host or reconfigure the SDD during the concurrent download of licensed machine code or you might lose your initial SDD configuration.

For information about performing concurrent download of LMC for virtualization products, see the Configuration Guide for your specific type and model.

Concurrent download of licensed machine code for DS3950, DS4000 and DS5000

If you are using the SDD multipath mode, you can concurrently download and install the licensed machine code while your applications continue to run, as long as you configure redundant paths to each storage controller port in addition to the multiple host adapter ports.

Because switching a device to another controller is a time-consuming recovery action and affects I/O performance, you can use this redundancy to avoid an

unnecessary controller failover if a path fails. Therefore, configure a minimum of four paths for each LUN with two host adapter ports and two storage controller ports where each host adapter port has redundancy to each storage controller port and vice versa.

Attention: Do not shut down the host or reconfigure the SDD during the concurrent download of licensed machine code or you might lose your initial SDD configuration.

Concurrent download of licensed machine code for IBM BladeCenter S SAS RAID Controller Module (RSSM)

Note: IBM does not support RSSM. For more information, contact the IBM Support team.

With the SDD multipath mode (configured with two paths per multipath device), you can concurrently download and install the licensed machine code while your applications continue to run.

During the code upgrade, each RSSM node is upgraded sequentially. The node that is being upgraded is temporarily unavailable, and all I/O operations to that node fail. However, failed I/O operations are directed to the other RSSM node, and applications do not see any I/O failures.

Attention: Do not shut down the host or reconfigure the SDD during the concurrent download of licensed machine code or you might lose your initial SDD configuration.

Active/Passive dual array controller path-selection algorithm for DS3950, DS4000 and DS5000 products

The DS4000 and DS5000 products are dual array controller disk subsystems. Each LUN is assigned to one controller, which is considered the owner, or the active controller, of a particular LUN. The other controller is considered as an alternate, or passive, controller. Thus, the SDD distinguishes the following paths to the DS4000 and DS5000 product LUN:

- Paths on the ownership (active) controller
- Paths on the alternate (passive) controller

With this type of active/passive dual-controller subsystem device, I/O can be sent only to the ownership controller. When the SDD selects paths for I/O, it selects paths that are connected only to the ownership controller. If there is no path on the ownership controller that can be used, SDD changes the LUN controller ownership to an alternate controller, switches the paths that were passive to active, and then selects these active paths for I/O.

Preferred node path-selection algorithm for the virtualization products

Virtualization products are two-controller disk subsystems. The SDD distinguishes the paths to a virtualization product LUN:

- Paths on the preferred controller
- Paths on the alternate controller

When the SDD selects paths for I/O, preference is always given to a path on the preferred controller. Therefore, in the selection algorithm, an initial attempt is made to select a path on the preferred controller. Only if no path can be used on the preferred controller is a path selected on the alternate controller. This means that the SDD automatically fails back to the preferred controller any time a path on the preferred controller becomes available during either manual or automatic recovery. Paths on the alternate controller are selected at random. If an error occurs and a path retry is required, retry paths are first selected on the preferred controller. If all retries fail on the preferred controller's paths, paths on the alternate controller are selected for retry. The following is the path selection algorithm for the SDD:

1. With all paths available, I/O is only routed to paths on the preferred controller.
2. If no path on the preferred controller is available, I/O failover processing to the alternate controller occurs.
3. When failover processing to the alternate controller has occurred, automatic I/O failback processing to the preferred controller occurs if a path on the preferred controller is available.

The following output of the **datapath query device** command shows that the preferred paths are being selected and shows the format of the virtualization product serial number.

```
DEV#: 0 DEVICE NAME: Disk0 Part0 TYPE: 2145 POLICY: OPTIMIZED
SERIAL: 6005676801800005F800000000000004
```

```
=====
```

Path#	Adapter/Hard Disk	State	Mode	Select	Errors
0	Scsi Port4 Bus0/Disk0 Part0	OPEN	NORMAL	501876	0
1	Scsi Port4 Bus0/Disk0 Part0	OPEN	NORMAL	501238	0
2	Scsi Port4 Bus0/Disk0 Part0	OPEN	NORMAL	0	0
3	Scsi Port4 Bus0/Disk0 Part0	OPEN	NORMAL	0	0
4	Scsi Port5 Bus0/Disk0 Part0	OPEN	NORMAL	499575	0
5	Scsi Port5 Bus0/Disk0 Part0	OPEN	NORMAL	500698	0
6	Scsi Port5 Bus0/Disk0 Part0	OPEN	NORMAL	0	0
7	Scsi Port5 Bus0/Disk0 Part0	OPEN	NORMAL	0	0

Open HyperSwap Replication

Open HyperSwap replication is a special Metro Mirror replication method that is designed to automatically failover I/O from the primary logical devices to the secondary logical devices in the event of a primary disk storage system failure. This function can be done with minimal disruption to the applications that are using the logical devices.

Open HyperSwap replication applies to both planned and unplanned replication sessions. When a session has Open HyperSwap enabled, an I/O error on the primary site automatically causes the I/O to switch to the secondary site without any user interaction, and with minimal application impact. In addition, while Open HyperSwap is enabled, the Metro Mirror session supports disaster recovery. If a write is successful on the primary site but is unable to get replicated on the secondary site, IBM Tivoli Storage Productivity Center for Replication suspends the entire set of data consistency checking, thus ensuring that a consistent copy of the data exists on the secondary site. If the system fails, this data might not be the latest data, but the data should be consistent and allow the user to manually switch host servers to the secondary site.

You can control Open HyperSwap from any system that is running IBM Tivoli Storage Productivity Center for Replication (AIX, Windows, Linux, or z/OS®).

However, the volumes that are involved with Open HyperSwap must be attached to an AIX system that is connected to IBM Tivoli Storage Productivity Center for Replication.

SDD distinguishes the paths of the source volume from the paths of the target volume on an Open HyperSwap copy set. With an Open HyperSwap device, I/O can only be sent to the source volume, so when SDD selects paths for I/O, it only selects paths that are connected to the source volume. If there is no path on the source volume that can be used, SDD will initiate the Open HyperSwap request to Tivoli Storage Productivity Center for Replication and work together to perform the swap. After the swap, SDD will select the target volume paths for I/O.

The following output of the `pcmpath query device` command shows that the target volume paths are being selected.

```
DEV#: 14 DEVICE NAME: hdisk14 TYPE: 2107900 ALGORITHM: Load Balance
SESSION NAME: session1
OS Direction: H1<-H2
=====
PRIMARY SERIAL: 25252520000
-----
Path#   Adapter/Path Name   State   Mode     Select   Errors
  0     fscsi0/path0        OPEN    NORMAL   6091     0
  1     fscsi0/path2        OPEN    NORMAL   6300     0
  2     fscsil/path4        OPEN    NORMAL   6294     0
  3     fscsil/path5        OPEN    NORMAL   6187     0

SECONDARY SERIAL: 34343430000 *
-----
Path#   Adapter/Path Name   State   Mode     Select   Errors
  4     fscsi0/path1        OPEN    NORMAL   59463    0
  5     fscsi0/path3        OPEN    NORMAL   59250    0
  6     fscsil/path6        OPEN    NORMAL   59258    0
  7     fscsil/path7        OPEN    NORMAL   59364    0
```

In the preceding example, the source volume is on site one, and the target volume is on site two. The output shows that after the swap from site one to site two, SDD selects paths of devices on site two.

Note: The primary serial is not always the source volume. The primary serial is the serial number of the volume on site one, and secondary serial is the serial number of the volume on site two.

Chapter 2. Using the SDD on an AIX host system

This chapter provides step-by-step procedures for installing, configuring, upgrading, and removing the SDD on an AIX host system that is attached to a supported storage device.

Beginning with SDD 1.4.0.5, the SDD supports the coexistence of ESS and SAN Volume Controller devices.

Beginning with SDD 1.6.0.0, the SDD supports the coexistence of all supported storage devices.

For updated and additional information that is not included in this chapter, see the Readme file on the CD-ROM or visit the SDD website:

www-947.ibm.com/support/entry/portal/product/system_storage/storage_software/other_software_products/system_storage_multipath_subsystem_device_driver?productContext=-1564959213

Supported SDD features

The following SDD features are supported in this release:

- 32- and 64-bit kernels
- Support for ESS, DS8000, DS6000 and virtualization products
- Preferred node path-selection algorithm for DS6000 and virtualization products
- Changing the SDD path-selection algorithm dynamically. Five path-selection algorithms are supported:
 - Failover
 - Round robin
 - Round robin sequential
 - Load balancing
 - Load balancing sequential
- Dynamically adding paths to the SDD vpath devices
- Dynamically opening an invalid or close_dead path
- Dynamically removing or replacing PCI adapters or paths
- Fibre-channel dynamic device tracking
- SDD server daemon support
- Support for HACMP™
- Support for secondary-system paging
- Support for load-balancing and failover protection for AIX applications and LVM
- SDD utility programs
- Support for SCSI-3 persistent reserve functions
- Support for AIX trace functions
- Support more than 512 SAN Volume Controller devices from multiple SAN Volume Controller clusters on an AIX host
- Storage I/O priority feature in DS6000 and DS8000, only with AIX53 TL04 or later and with 64-bit kernel

- Two types of reserve policies: No reserve and Persistent reserve exclusive host
- General Parallel File System (GPFS™)
- Virtual I/O Server with AIX 5.3 or later
- Dual Virtual I/O Server with AIX 5.3 or later

For more information about Virtual I/O Server, go to the following website:
www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/home.html

Verifying the hardware and software requirements

You must install the following hardware and software components to ensure that the SDD installs and operates successfully.

Hardware

The following hardware components are needed:

- One or more supported storage devices.
- A switch if using a SAN Volume Controller (no direct attachment allowed for SAN Volume Controller)
- Host system
- SCSI adapters and cables (for ESS only)
- Fibre-channel adapters and cables

Software

The following software components are needed:

- AIX operating system.
 Starting with SDD 1.6.1.0, the SDD package for AIX 5.3 (devices.sdd.53.rte) requires AIX53 TL04 with APAR IY76997.
 Starting with SDD 1.6.2.0, the SDD package for AIX 5.2 (devices.sdd.52.rte) requires AIX52 TL08 or later and the SDD package for AIX 5.3 (devices.sdd.53.rte) requires AIX53 TL04 or later.
- SCSI and fibre-channel device drivers
- ibm2105.rte package for ESS devices (devices.scsi.disk.ibm2105.rte or devices.fcp.disk.ibm2105.rte package if using NIM)
- devices.fcp.disk.ibm.rte for DS8000, DS6000, and SAN Volume Controller

Packages for SDD 1.4.0.0 (and later) will be using new package names in order to comply with AIX packaging rules and allow for NIM installation. Table 3 shows the package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0 (or later).

Table 3. Package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0 (or later)

SDD 1.3.3.x	SDD 1.4.0.0 (or later)	Notes®
ibmSdd_432.rte	Not applicable	Obsolete. This package has been merged with devices.sdd.43.rte.
ibmSdd_433.rte	devices.sdd.43.rte	Not applicable
ibmSdd_510.rte	Not applicable	Obsolete. This package has been merged with devices.sdd.51.rte.
ibmSdd_510nchacmp.rte	devices.sdd.51.rte	Not applicable
Not applicable	devices.sdd.52.rte	New package for AIX 5.2.0 (or later).

Table 3. Package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0 (or later) (continued)

SDD 1.3.3.x	SDD 1.4.0.0 (or later)	Notes [®]
Not applicable	devices.sdd.53.rte	New package for AIX 5.3.0 (or later).
Not applicable	devices.sdd.61.rte	New package for AIX 6.1.0 (or later).
<p>Note:</p> <ol style="list-style-type: none"> 1. SDD 1.4.0.0 (or later) no longer releases separate packages for concurrent and nonconcurrent High Availability Cluster Multiprocessing (HACMP). Both concurrent and nonconcurrent HACMP functions are now incorporated into one package for each AIX kernel level. 2. A persistent reserve issue arises when migrating from SDD to non-SDD volume groups after a reboot. This special case only occurs if the volume group was varied on prior to the reboot and auto varyon was not set when the volume group was created. See “Understanding the persistent reserve issue when migrating from SDD to non-SDD volume groups after a system reboot” on page 71 for more information. 		

Unsupported environments

The SDD does not support:

- A host system with both a SCSI and fibre-channel connection to a shared ESS logical unit number (LUN).
- Placing system primary paging devices (for example, /dev/hd6) on an SDD vpath device
- Any application that depends on a SCSI-2 reserve and release device on AIX
- Single-path mode during concurrent download of licensed machine code nor during any disk storage systems concurrent maintenance that impacts the path attachment, such as a disk storage systems host-bay-adapter replacement
- Multipathing to a system boot device
- Configuring the SDD vpath devices as system primary or secondary dump devices
- More than 600 SDD vpath devices if the host system is running AIX 4.3.3 or AIX 5.1.0
- More than 1200 SDD vpath devices if the host system is running AIX 5.2, AIX 5.3, or AIX 6.1
- DS8000, DS6000, and SAN Volume Controller with SCSI connectivity
- Multiple AIX servers without the SDD-supported clustering software, such as HACMP, installed

Host system requirements

To successfully install the SDD for supported storage devices, you must have AIX 4.3, AIX 5.1, AIX 5.2, AIX 5.3, or AIX 6.1 installed on your host system.

Starting with SDD 1.6.1.0, the SDD package for AIX 5.3 (devices.sdd.53.rte) requires AIX53 TL04 or later with APAR IY76997.

Starting with SDD 1.6.2.0, the SDD package for AIX 5.2 (devices.sdd.52.rte) requires AIX52 TL08 or later and the SDD package for AIX 5.3 (devices.sdd.53.rte) requires AIX53 TL04 or later.

You **must** check for and download the latest authorized program analysis reports (APARS), maintenance-level fixes, and microcode updates from the following website:

www-933.ibm.com/support/fixcentral/

Disk storage systems requirements

To successfully install the SDD:

- Ensure that the disk storage system devices are configured as:
 - For ESS:
 - IBM 2105xxx (SCSI-attached device)
where *xxx* represents the disk storage system model number.
 - IBM FC 2105 (fibre-channel-attached device)
 - For DS8000, IBM FC 2107
 - For DS6000, IBM FC 1750

Virtualization products requirements

To successfully install the SDD, ensure that the SAN Volume Controller devices are configured either as SAN Volume Controller Device or FC 2145.

SCSI requirements for ESS

To use the SDD SCSI support for ESS, ensure that your host system meets the following requirements:

- The bos.adt package is installed. The host system can be a single processor or a multiprocessor system, such as Symmetric Multi-Processor (SMP).
- A SCSI cable connects each SCSI host adapter to an ESS port.
- If you need the SDD input/output (I/O) load-balancing and failover features, ensure that a minimum of two SCSI adapters are installed.

Fibre requirements

You must check for and download the latest fibre-channel device driver APARS, maintenance-level fixes, and microcode updates from the following website:

www-933.ibm.com/support/fixcentral/

Note:

1. If your host has only one fibre-channel adapter, it requires you to connect through a switch to multiple disk storage system ports. You must have at least two fibre-channel adapters to prevent data loss due to adapter hardware failure or software failure.
2. The SAN Volume Controller *always* requires that the host be connected through a switch. For more information, see the *IBM System Storage SAN Volume Controller Model 2145–8A4 Hardware Installation Guide*.

For information about the fibre-channel adapters that can be used on your AIX host system, go to the following website:

www.ibm.com/servers/storage/support

To use the SDD fibre-channel support, ensure that your host system meets the following requirements:

- The AIX host system is an IBM RS/6000[®]™ or IBM System p[®]™ with AIX 4.3.3 (or later).

- The AIX host system has the fibre-channel device drivers installed along with all latest APARs.
- The bos.adt package is installed. The host system can be a single processor or a multiprocessor system, such as SMP.
- A fiber-optic cable connects each fibre-channel adapter to a disk storage system port.
- A fiber-optic cable connects each SAN Volume Controller fibre-channel adapter to a switch. The switch must also be configured correctly. See the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide* for information about the SAN Volume Controller.
- If you need the SDD I/O load-balancing and failover features, ensure that a minimum of two paths to a device are attached.

Preparing for the SDD installation

Before you install the SDD, you must complete the tasks identified in the following sections:

1. “Configuring the disk storage system”
2. “Configuring the virtualization products”
3. “Installing the AIX fibre-channel device drivers” on page 16
4. “Installing and upgrading the AIX SDD host attachment” on page 17
5. “Configuring fibre-channel-attached devices” on page 18
6. “Verifying the adapter firmware level” on page 19
7. “Determining if the sddServer for Expert is installed” on page 19
8. “Determining the installation package” on page 20
9. “Determining the installation type” on page 22

Note: The SDD allows the manual exclusion of supported devices from the SDD configuration. If you want to manually exclude supported devices (hdisk) from the SDD configuration, you must use the **excludesddcfg** command before configuring the SDD vpath devices. The **excludesddcfg** command reads the unique serial number of a device (hdisk) and saves the serial number in an exclude file. For detailed information about the **excludesddcfg** command, see “Manual exclusion of devices from the SDD configuration” on page 51.

Configuring the disk storage system

Before you install the SDD, you must configure:

- The disk storage system to your host system.
- A minimum of two independent paths that share the same logical unit. With a single-path configuration, failover protection is not provided.

For more information about how to configure your disk storage system, see the *Introduction and Planning Guide*.

Note: Ensure that the correct host attachment package for your storage device is installed on the host system.

Configuring the virtualization products

Before you install the SDD, you must configure:

- The virtualization product to your host system.

- SDD requires a minimum of two independent paths that share the same logical unit. With a single-path configuration, failover protection is not provided.

For more information about how to configure your SAN Volume Controller, see the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

Note: Ensure that the `devices.fcp.disk.ibm.rte` host attachment package is installed before configuring the virtualization product.

Installing the AIX fibre-channel device drivers

You must check for the latest information on fibre-channel device driver APARs, maintenance-level fixes, and microcode updates at the following website:

www-933.ibm.com/support/fixcentral/

Complete the following steps to install the AIX fibre-channel device drivers from the AIX compact disk:

1. Log in as the root user.
2. Load the compact disc into the CD-ROM drive.
3. From your desktop command prompt, enter **smitty install_update** and press **Enter** to go directly to the installation panels. The Install and Update Software menu is displayed.
4. Select **Install Software** and press **Enter**.
5. Press **F4** to display the INPUT Device/Directory for Software panel.
6. Select the compact disc drive that you are using for the installation; for example, `/dev/cd0`, and press **Enter**.
7. Press **Enter** again. The Install Software panel is displayed.
8. Select **Software to Install** and press **F4**. The Software to Install panel is displayed.
9. The fibre-channel device drivers include the following installation packages:

devices.pci.df1000f9

The adapter device driver for RS/6000 or IBM System p with feature code 6228.

devices.pci.df1000f7

The adapter device driver for RS/6000 or IBM System p with feature code 6227.

devices.common.IBM.fc

The FCP protocol driver.

devices.fcp.disk

The FCP disk driver.

devices.pci.df1080f9

The adapter device driver for RS/6000 or IBM System p with feature code 6239.

Select each one by highlighting it and pressing **F7**.

10. Press **Enter**. The Install and Update from LATEST Available Software panel is displayed with the name of the software you selected to install.
11. Check the default option settings to ensure that they are what you need.
12. Press **Enter** to install. SMIT responds with the following message:

```

+-----+
| ARE YOU SURE??                               |
| Continuing may delete information you may    |
| want to keep. 413                           |
| This is your last chance to stop before     |
| continuing. 415                             |
+-----+

```

13. Press **Enter** to continue. The installation process can take several minutes to complete.
14. When the installation is complete, press **F10** to exit from SMIT. Remove the compact disc.
15. Check to see if the correct APARs are installed by issuing the following command:

```
instfix -i | grep IYnnnnn
```

where *nnnnn* represents the APAR numbers.

If the APARs are listed, that means that they are installed. If they are installed, go to “Configuring fibre-channel-attached devices” on page 18. Otherwise, go to step 3.

16. Repeat steps 1 through 14 to install the APARs.

Uninstalling the AIX fibre-channel device drivers

The following steps describe how to uninstall the AIX fibre-channel device drivers. There are two methods for uninstalling all of your fibre-channel device drivers:

- **smitty deinstall** command
- **installp** command

Using the smitty deinstall command

Complete the following steps to use the **smitty deinstall** command:

1. Enter **smitty deinstall** at the AIX command prompt and press **Enter**. The Remove Installed Software panel is displayed.
2. Press **F4**. All of the software that is installed is displayed.
3. Select the file name of the fibre-channel device driver that you want to uninstall. Press **Enter**. The selected file name is displayed in the **Software Name** field of the Remove Installed Software panel.
4. Use the Tab key to toggle to **No** in the **PREVIEW Only?** field. Press **Enter**. The uninstallation process begins.

Using the installp command

Complete the following steps to use the **installp** command from the AIX command line:

1. Enter **installp -ug devices.pci.df1000f9** and press **Enter**.
2. Enter **installp -ug devices.pci.df1000f7** and press **Enter**.
3. Enter **installp -ug devices.pci.df1080f9** and press **Enter**.
4. Enter **installp -ug devices.common.IBM.fc** and press **Enter**.
5. Enter **installp -ug devices.fcp.disk** and press **Enter**.

Installing and upgrading the AIX SDD host attachment

The host attachments for SDD add 2105, 2145, 2107, or 1750 device information to allow AIX to properly configure 2105, 2145, 2107, or 1750 hdisks. This device information allows AIX to perform the following actions:

- Identify the hdisks as 2105, 2145, 2107, or a 1750 hdisks.
- Set default hdisk attributes, such as `queue_depth` and `timeout` values.

- Indicate to the AIX device driver configure method to configure the hdisk as a non-MPIO-capable device.

The following names are AIX SDD Host Attachment packages:

- `ibm2105.rte.tar` for 2105 devices. Also download the SCSI (`devices.scsi.disk.ibm2105.rte.tar`) or FC (`devices.fcp.disk.ibm2105.rte.tar`) ESS NIM support add on if you are using NIM.
- `devices.fcp.disk.ibm.rte.tar` for 2145, 2107, and 1750 devices. A NIM support add on is not necessary for the SAN Volume Controller, DS8000, or DS6000 Host Attachment.

Note: SDD 1.6.2.0 (or later) on AIX 5.3 or SDD 1.7.0.0 (or later) on AIX 5.3 (or later) requires one or both of the following packages:

- `ibm2105.rte 32.6.100.27` (or later)
- `devices.fcp.disk.ibm.rte 1.0.0.6` (or later)

See the SDD Readme file for the latest Host Attachment level for the SDD version that you are installing.

For AIX host attachment installation and upgrade instructions, see the *Host Systems Attachment Guide* for your storage system.

Configuring fibre-channel-attached devices

The newly installed fibre-channel-attached devices must be configured before you can use them. Use one of the following commands to configure these devices:

- **cfgmgr** command

Note: In AIX 4.3 or 5.1, you might need to issue the **cfgmgr** command once for each host adapter each time you add a new device.

After the command prompt appears, use the **lsdev -Cc disk** command to check the Fibre Channel Protocol (FCP) disk configuration. If the FCP devices are configured correctly, they are in the *Available* state. If the FCP devices are configured correctly, go to “Verifying the adapter firmware level” on page 19 to determine if the proper firmware level is installed.

- **shutdown -rF** command to restart the system.

After the system restarts, use the **lsdev -Cc disk** command to check the Fibre Channel Protocol (FCP) disk configuration. If the FCP devices are configured correctly, they are in the *Available* state. If the FCP devices are configured correctly, go to “Verifying the adapter firmware level” on page 19 to determine if the proper firmware level is installed.

Removing fibre-channel-attached devices

To remove all fibre-channel-attached devices, you must enter the following command for each installed FCP adapter:

```
rmdev -d1 fcsN -R
```

where *N* is the FCP adapter number. For example, if you have two installed FCP adapters (adapter 0 and adapter 1), you must enter both of the following commands:

```
rmdev -d1 fcs0 -R
rmdev -d1 fcs1 -R
```

Verifying the adapter firmware level

You must verify that your current adapter firmware is at the latest level. If your current adapter firmware is not at the latest level, you must upgrade to a new adapter firmware (microcode). To check the current supported firmware level for fibre-channel adapters, go to the following website:

www-933.ibm.com/support/fixcentral/

Tip:

- The current firmware level for the FC 6227 adapter is 3.30X1
- The current firmware level for the FC 6228 adapter is 3.93A0
- The current firmware level for the FC 6239 adapter is 1.92a1

Complete the following steps to verify the firmware level that is currently installed:

1. Enter the `lscfg -vl fcsN` command. The vital product data for the adapter is displayed.
2. Look at the **ZB** field. The **ZB** field should look similar to:

```
Device Specific.(ZB).....S2F3.30X1
```

To verify the firmware level, ignore the first three characters in the ZB field. In the example, the firmware level is 3.22A1

3. If the adapter firmware level is at the latest level, there is no need to upgrade; otherwise, the firmware level must be upgraded. For instructions on upgrading the firmware level, see the description for each firmware at:

www-933.ibm.com/support/fixcentral/

Determining if the `sddServer` for Expert is installed

If you previously installed the *stand-alone* version of the `sddServer` for IBM TotalStorage Expert V2R1 (ESS Expert) on your AIX host system, you must remove this stand-alone version of `sddServer` before you proceed with SDD 1.3.3.9 (or later) installation. The installation package for SDD 1.3.3.9 (or later) includes the SDD server daemon (also referred to as *sddsrv*), which incorporates the functionality of the stand-alone version of `sddServer` (for ESS Expert).

To determine if the stand-alone version of `sddServer` is installed on your host system, enter:

`lslpp -l sddServer.rte`

If you previously installed the `sddServer.rte` package, the output from the `lslpp -l sddServer.rte` command looks similar to this:

Fileset	Level	State	Description
Path: /usr/lib/objrepos sddServer.rte	1.0.0.0	COMMITTED	IBM SDD Server for AIX
Path: /etc/objrepos sddServer.rte	1.0.0.0	COMMITTED	IBM SDD Server for AIX

For more information about the SDD server daemon, go to “SDD server daemon” on page 67.

Understanding SDD support for IBM System p with static LPARs configured

The IBM System p server supports static LPARs as a standard feature, and users can partition them if they choose to do so. Each LPAR is composed of one or more processors, some dedicated memory, and dedicated I/O adapters. Each partition has an instance of an operating system and does not share IBM System p hardware resources with any other partition. So each partition functions the same way that it does on a stand-alone system. Storage subsystems need to be shared the same way that they have always been shared (shared storage pool, shared ports into the storage subsystem, and shared data on concurrent mode) where the application is capable of sharing data.

If a partition has multiple fibre-channel adapters that can see the same LUNs in a supported storage device, the path optimization can be performed on those adapters in the same way as in a stand-alone system. When the adapters are not shared with any other partitions, SCSI reservation, persistent reserve, and LUN level masking operate as expected (by being "bound" to an instance of the operating system).

The SDD provides the same functions on one of the partitions or LPARs of a IBM System p server as it does on a stand-alone server

Determining the installation package

Before you install the SDD on your AIX host system (4.3.3 or later), you need to determine the installation package that is appropriate for your AIX OS level. See Table 4 to determine the correct installation package.

Installation packages for 32-bit and 64-bit applications on AIX 4.3.3 (or later) host systems

Table 4. SDD 1.4.0.0 (or later) installation packages for different AIX OS levels and the supported AIX kernel mode, application mode, and interface

SDD installation package names	AIX OS level	AIX kernel mode	Application mode	SDD interface
devices.sdd.43.rte	AIX 4.3.3 ¹	32-bit	32-bit, 64-bit	LVM, raw device
devices.sdd.51.rte	AIX 5.1.0	32-bit, 64-bit	32-bit, 64-bit	LVM, raw device
devices.sdd.52.rte	AIX 5.2.0	32-bit, 64-bit	32-bit, 64-bit	LVM, raw device
devices.sdd.53.rte	AIX 5.3.0	32-bit, 64-bit	32-bit, 64-bit	LVM, raw device
devices.sdd.61.rte	AIX 6.1.0 ²	64-bit	32-bit, 64-bit	LVM, raw device

Note:

1. devices.sdd.43.rte is supported only by the ESS and virtualization products.
2. devices.sdd.61.rte supports only the 64-bit kernel.

Switching between 32-bit and 64-bit modes on AIX 5.1.0, AIX 5.2.0, and AIX 5.3.0 host systems

SDD supports AIX 5.1.0, AIX 5.2.0 and AIX 5.3.0 host systems that run in both 32-bit and 64-bit kernel modes. You can use the **bootinfo -K** or **ls -al /unix** command to check the current kernel mode in which your AIX 5.1.0, 5.2.0, or 5.3.0 host system is running.

The **bootinfo -K** command directly returns the kernel mode information of your host system. The **ls -al /unix** command displays the /unix link information. If the /unix links to /usr/lib/boot/unix_mp, your AIX host system runs in 32-bit mode. If the /unix links to /usr/lib/boot/unix_64, your AIX host system runs in 64-bit mode.

If your host system is currently running in 32-bit mode, you can switch it to 64-bit mode by typing the following commands in the given order:

```
ln -sf /usr/lib/boot/unix_64 /unix
ln -sf /usr/lib/boot/unix_64 /usr/lib/boot/unix
bosboot -ak /usr/lib/boot/unix_64
shutdown -Fr
```

The kernel mode of your AIX host system is switched to 64-bit mode after the system restarts.

If your host system is currently running in 64-bit mode, you can switch it to 32-bit mode by typing the following commands in the given order:

```
ln -sf /usr/lib/boot/unix_mp /unix
ln -sf /usr/lib/boot/unix_mp /usr/lib/boot/unix
bosboot -ak /usr/lib/boot/unix_mp
shutdown -Fr
```

The kernel mode of your AIX host system is switched to 32-bit mode after the system restarts.

Installation of major files on your AIX host system

The SDD installation package installs a number of major files on your AIX system. Table 5 lists the major files that are part of the SDD installation package.

Table 5. Major files included in the SDD installation package

File name	Description
defdpo	Define method of the SDD pseudo-parent data path optimizer (dpo).
cfgdpo	Configure method of the SDD pseudo-parent dpo.
define_vp	Define method of the SDD vpath devices.
addpaths	The command that dynamically adds more paths to SDD vpath devices while they are in <i>Available</i> state.
cfgvpath	Configure method of the SDD vpath devices.
chgvpath	Method to change vpath attributes.
cfallvpath	Fast-path configuration method to configure the SDD pseudo-parent dpo and all SDD vpath devices.
vpathdd	The SDD device driver.
hd2vp	The SDD script that converts an hdisk device volume group to an SDD vpath device volume group.

Table 5. Major files included in the SDD installation package (continued)

vp2hd	The SDD script that converts an SDD vpath device volume group to an hdisk device volume group.
datapath	The SDD driver console command tool.
lquerypr	The SDD driver persistent reserve command tool.
lsvpcfg	The SDD driver query configuration state command.
excludesddcfg	The SDD driver tool to exclude user-specified hdisk devices from the SDD vpath configuration.
mkvg4vp	The command that creates an SDD volume group.
extendvg4vp	The command that extends the SDD vpath devices to an SDD volume group.
dpovgfix	The command that fixes an SDD volume group that has mixed vpath and hdisk physical volumes.
savevg4vp	The command that backs up all files belonging to a specified volume group with the SDD vpath devices.
restvg4vp	The command that restores all files belonging to a specified volume group with the SDD vpath devices.
sddsrv	The SDD server daemon for path reclamation and probe.
sample_sddsrv.conf	The sample SDD server configuration file.
lvmrecover	The SDD script that restores a system's SDD vpath devices and LVM configuration when a migration failure occurs.
sddfcmap	The SDD tool that collects information on ESS SCSI or disk storage systems fibre-channel devices through SCSI commands.
sddgetdata	The SDD data collection tool for problem analysis.

Determining the installation type

Before you install the SDD on your AIX host system 4.3.3 (or later), you need to determine the installation type that is appropriate for your environment.

If there is no previous version of the SDD installed on the host system, see “Installing the SDD” on page 23 for instructions on installing and configuring the SDD. If there is a previous version of the SDD installed on the host system and you want to upgrade to one of the following packages:

- devices.sdd.43.rte
- devices.sdd.51.rte
- devices.sdd.52.rte
- devices.sdd.53.rte
- devices.sdd.61.rte

See “Upgrading the SDD packages automatically without system restart” on page 24 for instructions on upgrading the SDD.

If SDD 1.4.0.0 (or later) is installed on the host system and you have an SDD PTF that you want to apply to the system, see “Updating SDD packages by applying a program temporary fix” on page 29 for instructions. A PTF file has a file extension of bff (for example, devices.sdd.43.rte.2.1.0.1.bff) and requires special consideration when being installed.

Installing and upgrading the SDD

This section provides information about installing, upgrading, and verifying the SDD.

Installing the SDD

The SDD is released as an installation image. To install the SDD, use the installation package that is appropriate for your AIX OS level.

Table 4 on page 20 lists the correct installation packages for different AIX OS level and the supported AIX kernel mode, application mode, and interface.

You must have root access and AIX system administrator knowledge to install SDD.

If you are installing an older version of SDD when a newer version is already installed, you must first remove the newer version from your host system before you can install the older version of SDD.

Note: The following procedures assume that SDD will be used to access all of your single-path and multipath devices.

Use the System Management Interface Tool (SMIT) facility to install SDD. The SMIT facility has two interfaces, nongraphical and graphical. Enter **smitty** to invoke the nongraphical user interface or enter **smit** to invoke the graphical user interface (GUI).

Installing the SDD from CD-ROM

Throughout this SMIT procedure, `/dev/cd0` is used for the compact disc drive address. The drive address can be different in your environment. Complete the following SMIT steps to install the SDD package on your system.

1. Log in as the root user.
2. Load the compact disc into the CD-ROM drive.
3. From your desktop window, enter **smitty install_update** and press **Enter** to go directly to the installation panels. The Install and Update Software menu is displayed.
4. Select **Install Software** and press **Enter**.
5. Press **F4** to display the INPUT Device/Directory for Software panel.
6. Select the compact disc drive that you are using for the installation, for example, `/dev/cd0`; and press **Enter**.
7. Press **Enter** again. The Install Software panel is displayed.
8. Select **Software to Install** and press **F4**. The Software to Install panel is displayed.
9. Select the installation package that is appropriate for your environment.
10. Press **Enter**. The Install and Update from LATEST Available Software panel is displayed with the name of the software that you selected to install.
11. Check the default option settings to ensure that they are what you need.
12. Press **Enter** to install. SMIT responds with the following message:

```
ARE YOU SURE??  
Continuing may delete information you may want to keep.  
This is your last chance to stop before continuing.
```

13. Press **Enter** to continue. The installation process can take several minutes to complete.
14. When the installation is complete, press **F10** to exit from SMIT. Remove the compact disc.

Note: You do not need to reboot SDD even though the bosboot message indicates that a reboot is necessary.

Installing SDD from downloaded code

Use this procedure to install SDD from downloaded code:

1. Download code from the website that is appropriate for your system.
2. Move the downloaded files to the /usr/sys/inst.images directory.
3. Change directory to the /usr/sys/inst.images directory.
4. Untar the downloaded file: for example, **tar -xvf devices.sdd.52.rte.tar**
5. Update the .toc file. For example:

```
pwd
rm -i .toc
inutoc .
grep -i sdd .toc
```

This command reflects the newer SDD code version that will be updated.

6. Continue the installation by following the instructions beginning in step 3 on page 23.

Upgrading the SDD

Upgrading the SDD packages automatically without system restart

With SDD 1.4.0.0 (or later), a new feature is provided to migrate or upgrade SDD packages. This feature supports backup, restoration, and recovery of LVM configurations and SDD device configurations automatically on the server, as well as migration from non-PR to PR SDD packages. This is especially beneficial in a complex SAN environment where a system administrator has to maintain a large number of servers. During SDD migration or upgrade, the LVM and SDD device configuration of the host will automatically be removed, the new SDD package will be installed, and then the SDD device and LVM configuration of the host will be restored. This feature will support the following scenarios:

1. Package migration from a nonpersistent reserve package with version 1.3.1.3 (or later) to a persistent reserve package with version 1.4.0.0 (or later). That is, `ibmSdd_432.rte` → `devices.sdd.43.rte` and `ibmSdd_510.rte` → `devices.sdd.51.rte`.
2. Package migration from version 1.3.1.3 or later to version 1.4.0.0 or later. Migration from SDD version earlier than 1.3.1.3 is not supported.
3. Package upgrade from version 1.4.0.0 to a later version.

If the SDD currently installed on your host system is listed in Table 6, you can use this automatic migration to upgrade the SDD. If the SDD currently installed on your host system is not listed in Table 6, you will need to upgrade the SDD manually

Table 6. List of previously installed installation packages that are supported with the installation upgrade

Installation package name
ibmSdd_432.rte

Table 6. List of previously installed installation packages that are supported with the installation upgrade (continued)

Installation package name
ibmSdd.rte.432
ibmSdd_433.rte
ibmSdd.rte.433
ibmSdd_510.rte
ibmSdd_510nchacmp.rte
devices.sdd.43.rte
devices.sdd.51.rte
devices.sdd.52.rte
devices.sdd.53.rte
devices.sdd.61.rte

Beginning with SDD 1.6.0.0, SDD introduces a new feature in the configuration method to read the pvid from the physical disks and convert the pvid from hdisks to vpaths during the SDD vpath configuration. With this feature, you can skip the process of converting the pvid from hdisks to vpaths after configuring SDD devices. Furthermore, the SDD migration scripts can now skip the pvid conversion scripts. This tremendously reduces the SDD migration time, especially with a large number of SDD devices and LVM configuration environment.

Furthermore, the SDD now introduces two new environment variables that can be used in some configuration environments to customize the SDD migration and further reduce the time needed to migrate or upgrade the SDD. See “Customizing the SDD migration or upgrade” for details.

During the migration or upgrade of the SDD, the LVM configuration of the host will be removed, the new SDD package will be installed, and then the original LVM configuration of the host will be restored.

Preconditions for migration or upgrade:

The following are the preconditions for running the migration:

1. If HACMP is running, gracefully stop the cluster services.
2. If sddServer.rte (stand-alone IBM TotalStorage Expert SDD Server) is installed, uninstall sddServer.rte.
3. If there is any I/O running to SDD devices, stop these I/O activities.
4. Stop any activity related to system configuration changes. These activities are not allowed during the SDD migration or upgrade (for example, configuring more devices).
5. If there is active paging space created with the SDD devices, deactivate the paging space.

If any of the preceding preconditions are not met, the migration or upgrade will fail.

Customizing the SDD migration or upgrade: Beginning with SDD 1.6.0.0, SDD offers two new environment variables, SKIP_SDD_MIGRATION and SDDVG_NOT_RESERVED, for you to customize the SDD migration or upgrade to maximize performance. You can set these two variables based on the configuration

of your system. The following discussion explains the conditions and procedures for using these two environment variables.

SKIP_SDD_MIGRATION:

The `SKIP_SDD_MIGRATION` environment variable is an option available to bypass of the SDD automated migration process (backup, restoration, and recovery of LVM configurations and SDD device configurations). This variable could help to decrease SDD upgrade time if you choose to reboot the system after upgrading SDD. For example, you might choose this option if you are upgrading other software that requires a reboot on the host at the same time. Another example is if you have a large number of SDD devices and LVM configuration, and a system reboot is acceptable. In these cases, you might want to choose this option to skip the SDD automated migration process. If you choose to skip the SDD automated migration process, follow these procedures to perform an SDD upgrade:

1. Issue **export SKIP_SDD_MIGRATION=YES** to set the `SKIP_SDD_MIGRATION` environment variable.
2. Issue **smitty install** to install SDD.
3. Reboot the system.
4. Issue **varyonvg vg_name** for the volume groups that are not auto-varied on after reboot.
5. Issue **mount filesystem-name** to mount the file system.

SDDVG_NOT_RESERVED:

`SDDVG_NOT_RESERVED` is an environment variable to indicate to the SDD migration script whether the host has any SDD volume group reserved by another host. If the host has any SDD volume group reserved by another host, set this variable to *NO*. Otherwise, you should set this variable to *YES*. If this variable is not set, the SDD migration script will assume the value to be *NO*.

When this variable is set to *YES*, the SDD migration script skips some procedures. This dramatically reduces the SDD migration time. If `SDDVG_NOT_RESERVED` is set to *NO*, the SDD migration script makes certain assumptions and runs more steps.

Set this variable to *YES* if the host is:

1. A completely stand-alone host, that is, not sharing LUNs with any other host
2. A host in a clustering environment but all the volume groups (including the volume groups that belong to a cluster software resource group) are configured for concurrent access only
3. A host in a clustering environment with nonconcurrent volume groups but all the nonconcurrent volume groups on all the hosts are varied off. That is, no other node has made reserve on the SDD volume groups.

If the host does not meet the any of these three conditions, set `SDDVG_NOT_RESERVED` to *NO*, so that the SDD migration script runs the **vp2hd** pvid conversion script to save the pvid under hdisks. Follow these procedures to perform SDD migration with this variable:

1. Issue **export SDDVG_NOT_RESERVED=NO** or **export SDDVG_NOT_RESERVED=YES** to set the `SDDVG_NOT_RESERVED` environment variable
2. Follow the procedures in "Procedures for automatic migration or upgrade."

Procedures for automatic migration or upgrade:

To start the SDD migration or upgrade:

1. Install the new SDD package by entering the **smitty install** command. The migration or upgrade scripts run as part of the installation procedure that is initiated by the **smitty install** command. These scripts save SDD-related LVM configuration on the system.

SDD does not support mixed volume groups with the SDD vpath devices and supported storage hdisk devices. A volume group contains the SDD vpath devices or only supported storage hdisk devices. If you do have a mixed volume group, the SDD migration or upgrade script fixes it by changing only the volume group to contain the SDD vpath devices. The following message displays when the SDD migration or upgrade script fixes the mixed volume group:

```
<volume group> has a mixed of SDD and non-SDD devices.  
dpovgfix <volume group> is run to correct it.  
Mixed volume group <volume group> is converted to SDD devices successfully!
```

The following messages indicate that the preuninstallation operations of the SDD are successful:

```
LVM configuration is saved successfully.  
All mounted file systems are unmounted.  
All varied-on volume groups are varied off.  
All volume groups created on SDD devices are converted to non-SDD devices.  
SDD Server is stopped.  
All SDD devices are removed.  
Ready for deinstallation of SDD!
```

2. The older SDD is uninstalled before new SDD will be installed.
3. The migration or upgrade script automatically configures the SDD devices and restores the original LVM configuration.

The following messages indicate that the postinstallation of SDD is successful:
Original lvm configuration is restored successfully!

Error recovery for migration or upgrade: If any error occurred during the preinstallation or postinstallation procedures, such as disconnection of cables, you can recover the migration or upgrade. There are two common ways that the migration or the upgrade can fail:

Case 1: Smitty install failed.

Smitty install fails if there is an error during the preuninstallation activities for the older SDD package. An error message indicating the error is printed, so you can identify and fix the problem. Use the **smitty install** command to install the new SDD package again.

Case 2: Smitty install exits with an OK prompt but configuration of SDD devices or LVM restoration failed.

If there is an error during the postinstallation (either the configuration of SDD devices has failed or LVM restoration has failed), the new SDD package is still successfully installed. Thus, the **Smitty install** exits with an OK prompt. However, an error message indicating the error is printed, so you can identify and fix the problem. Then, run the shell script **lvmrecover** to configure SDD devices and automatically recover the original LVM configuration.

Upgrading SDD manually

The following section describes the procedure to migrate or upgrade SDD manually. See “Upgrading the SDD packages automatically without system restart” on page 24 for information about migrating or upgrading SDD automatically.

A manual migration or upgrade is required if you are upgrading from a previous version of the SDD package not listed in Table 6 on page 24. You must uninstall the existing SDD and then manually install the new version of SDD in these cases.

Complete the following steps to upgrade SDD:

1. Enter the **lspv** command to find out all the SDD volume groups.
2. Enter the **lsvgfs** command for each SDD volume group to find out which file systems are mounted. Enter the following command:
lsvgfs vg_name
3. Enter the **umount** command to unmount all file systems belonging to the SDD volume groups. Enter the following command:
umount filesystem_name
4. Enter the **varyoffvg** command to vary off the volume groups. Enter the following command:
varyoffvg vg_name
5. If you are upgrading to an SDD version earlier than 1.6.0.0; or if you are upgrading to SDD 1.6.0.0 or later and your host is in a HACMP environment with nonconcurrent volume groups that are varied-on on other host, that is, reserved by other host, run the **vp2hd volume_group_name** script to convert the volume group from the SDD vpath devices to supported storage hdisk devices. Otherwise, you skip this step.
6. Stop the SDD server by entering the following command:
stopsrc -s sddsrv
7. Remove all the SDD vpath devices. Enter the following command:
rmdev -dl dpo -R
8. Use the **smitty** command to uninstall the SDD. Enter **smitty deinstall** and press **Enter**. The uninstallation process begins. Complete the uninstallation process. See “Removing SDD from an AIX host system” on page 49 for the step-by-step procedure for uninstalling the SDD.
9. If you need to upgrade the AIX operating system, for example, from AIX 4.3 to AIX 5.1, you could perform the upgrade now. If required, reboot the system after the operating system upgrade.
10. Use the **smitty** command to install the newer version of the SDD from the compact disc. Enter **smitty install** and press **Enter**. The installation process begins. Go to “Installing and upgrading the SDD” on page 23 to complete the installation process.
11. Use the **smitty device** command to configure all the SDD vpath devices to the *Available* state. See “Configuring SDD” on page 44 for a step-by-step procedure for configuring devices.
12. Enter the **lsvpcfg** command to verify the SDD configuration. Enter the following command:
lsvpcfg
13. If you are upgrading to an SDD version earlier than 1.6.0.0, run the **hd2vp volume_group_name** script for each SDD volume group to convert the physical volumes from supported storage hdisk devices back to the SDD vpath devices. Enter the following command:

`hd2vp volume_group_name`

14. Enter the **varyonvg** command for each volume group that was previously varied offline. Enter the following command:
`varyonvg vg_name`
15. Enter the **lspv** command to verify that all physical volumes of the SDD volume groups are SDD vpath devices.
16. Enter the **mount** command to mount all file systems that were unmounted in step 3 on page 28. Enter the following command:
`mount filesystem-name`

Attention: If the physical volumes on an SDD volume group's physical volumes are mixed with hdisk devices and SDD vpath devices, you must run the **dpovgfix** utility to fix this problem. Otherwise, SDD will not function properly. Enter the **dpovgfix vg_name** command to fix this problem.

Updating SDD packages by applying a program temporary fix

SDD 1.4.0.0 and later allows users to update SDD by installing a program temporary fix (PTF). A PTF file has a file extension of *bff* (for example, `devices.sdd.43.rte.2.1.0.1.bff`) and can either be applied or committed when it is installed. If the PTF is committed, the update to SDD is permanent; to remove the PTF, you must uninstall SDD. If the PTF is applied, you can choose to commit or to reject the PTF at a later time. If you decide to reject the PTF, you will not need to uninstall SDD from the host system.

Use the System Management Interface Tool (SMIT) facility to update SDD.

Throughout this SMIT procedure, `/dev/cd0` is used for the compact disc drive address. The drive address can be different in your environment.

Complete the following SMIT steps to update the SDD package on your system:

1. Log in as the root user.
2. Load the compact disc into the CD-ROM drive.
3. From your desktop window, enter **smitty install_update** and press **Enter** to go directly to the installation panels. The Install and Update Software menu is displayed.
4. Select **Install Software** and press **Enter**.
5. Press **F4** to display the INPUT Device/Directory for Software panel.
6. Select the compact disc drive that you are using for the installation (for example, `/dev/cd0`) and press **Enter**.
7. Press **Enter** again. The Install Software panel is displayed.
8. Select **Software to Install** and press **F4**. The Software to Install panel is displayed.
9. Select the PTF package that you want to install.
10. Press **Enter**. The Install and Update from LATEST Available Software panel is displayed with the name of the software that you selected to install.
11. If you only want to apply the PTF, select **Commit software Updates?** and tab to change the entry to *no*. The default setting is to commit the PTF. If you specify *no* to **Commit Software Updates?**, be sure that you specify *yes* to **Save Replaced Files?**.
12. Check the other default option settings to ensure that they are what you need.
13. Press **Enter** to install. SMIT responds with the following message:

```

+-----+
| ARE YOU SURE??                               |
| Continuing may delete information you may    |
| want to keep.                               |
| This is your last chance to stop before     |
| continuing.                                 |
+-----+

```

14. Press **Enter** to continue. The installation process can take several minutes to complete.
15. When the installation is complete, press **F10** to exit from SMIT.
16. Remove the compact disc.

Note: You do not need to reboot SDD even though the bosboot message indicates that a reboot is necessary.

Committing or Rejecting a PTF Update:

Before you reject a PTF update, you need to stop sddsrv and remove all SDD devices. The following steps will guide you through this process. If you want to commit a package, you will not need to complete these steps. Follow these steps prior to rejecting a PTF update:

1. Stop SDD Server. Enter the following command:
stopsrc -s sddsrv
2. Enter the **lspv** command to find out all the SDD volume groups.
3. Enter the **lsvgfs** command for each SDD volume group to find out which file systems are mounted. Enter the following command:
lsvgfs vg_name
4. Enter the **umount** command to unmount all file systems belonging to SDD volume groups. Enter the following command:
umount filesystem_name
5. Enter the **varyoffvg** command to vary off the volume groups. Enter the following command:
varyoffvg vg_name
6. If you are downgrading to an SDD version earlier than 1.6.0.0 or if you are downgrading to SDD 1.6.0.0 or later but your host is in a HACMP environment with nonconcurrent volume groups that are varied-on on other host (that is, reserved by other host), run the **vp2hd volume_group_name** script to convert the volume group from SDD vpath devices to supported storage hdisk devices. Otherwise, you skip this step.
7. Remove all SDD devices. Enter the following command:
rmdev -dl dpo -R

Complete the following steps to commit or reject a PTF update with the SMIT facility.

1. Log in as the root user.
2. From your desktop window, enter **smitty install** and press **Enter** to go directly to the installation panels. The Software Installation and Maintenance menu is displayed.
3. Select **Software Maintenance and Utilities** and press **Enter**.
4. Select **Commit Applied Software Updates** to commit the PTF or select **Reject Applied Software Updates** to reject the PTF.
5. Press **Enter**. The Commit Applied Software Updates panel is displayed or the Reject Applied Software Updates panel is displayed.
6. Select **Software name** and press **F4**. The software name panel is displayed.
7. Select the Software package that you want to commit or reject.

8. Check the default option settings to ensure that they are what you need.
9. Press **Enter**. SMIT responds with the following message:

```

+-----+
| ARE YOU SURE??                               |
| Continuing may delete information you may    |
| want to keep.                               |
| This is your last chance to stop before     |
| continuing.                                 |
+-----+

```

10. Press **Enter** to continue. The commit or reject process can take several minutes to complete.
11. When the installation is complete, press **F10** to exit from SMIT.

Note: You do not need to reboot SDD even though the bosboot message might indicate that a reboot is necessary.

After the procedure to reject a PTF update completes successfully:

1. Use the **smitty device** command to configure all the SDD vpath devices to the *Available* state. See “Configuring fibre-channel-attached devices” on page 18 for a step-by-step procedure for configuring devices.
2. Enter the **lsvpcfg** command to verify the SDD configuration. Enter the following command:


```
lsvpcfg
```
3. If you have downgraded to an SDD version earlier than 1.6.0.0, run the **hd2vp** script for each SDD volume group to convert the physical volumes from supported storage hdisk devices back to SDD vpath devices. Enter the following command:


```
hd2vp vg_name
```
4. Enter the **varyonvg** command for each volume group that was previously varied offline. Enter the following command:


```
varyonvg vg_name
```
5. Enter the **lspv** command to verify that all physical volumes of the SDD volume groups are SDD vpath devices.
6. Enter the **mount** command to mount all file systems that were unmounted in step 4. Enter the following command:


```
mount filesystem-name
```

Note: If the physical volumes on an SDD volume group's 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. Enter the **dpovgfix vg_name** command to fix this problem.

7. Start the SDD Server. Enter the following command:


```
startsrc -s sddsrv
```

Upgrading AIX OS or host attachment and SDD packages

SDD provides different packages to match the AIX OS level. If an AIX system is going to be upgraded to a different OS level, you need to install the corresponding SDD package for that OS level. Automatic migration of an SDD package from an earlier OS level to a later OS level after an OS upgrade is not supported. For example, automatic migration from devices.sdd.43.rte to devices.sdd.51.rte after an OS upgrade from AIX 4.3 to AIX 5.1, or automatic migration from devices.sdd.51.rte to devices.sdd.52.rte after an OS upgrade from AIX 5.1 to 5.2, is not supported.

Use the procedures in either Case 1 or Case 2 when you want to upgrade one of the following systems:

- AIX OS only*
- Host attachment and AIX OS*
- SDD and AIX OS
- Host attachment and SDD
- Host attachment only
- SDD, host attachment, and AIX OS

* Upgrading the AIX OS always requires you to install the SDD that corresponds to the new AIX OS level.

Before you upgrade, ensure that *rootvg* is either on local SCSI disks or a SAN boot disk.

If you want to upgrade SDD only, see “Upgrading the SDD packages automatically without system restart” on page 24 or “Upgrading SDD manually” on page 28.

Case 1: In stand-alone host or with the HACMP services stopped

If the system is a stand-alone host or if you can stop the HACMP cluster services, these are the procedures for upgrading the software.

1. Stop all activities related to SDD devices:
 - a. Stop applications running on SDD volume groups or file systems.
 - b. If your host is in an HACMP environment, stop the cluster services in an orderly manner.
 - c. If you are not upgrading AIX OS and you can reboot the system, go to Step 5 on page 33. Otherwise, continue to the next step.
 - d. If there is active paging space created with SDD devices, deactivate the paging space.
 - e. Use the **umount** command to unmount all file systems of SDD volume groups.
 - f. Vary off all SDD volume groups.
2. Remove SDD vpath devices using the **rmdev -dl dpo -R** command. See “Special considerations when you uninstall SDD in the NIM environment” on page 50.
3. If *rootvg* is on local SCSI disks, follow these procedures. Otherwise, if *rootvg* is on a SAN boot disk, skip to step 4.
 - a. Remove hdisk devices using the following command:

lsdev -C -t 2105* -F name		xargs -n1 rmdev -dl	for 2105 devices
lsdev -C -t 2145* -F name		xargs -n1 rmdev -dl	for 2145 devices
lsdev -C -t 2107* -F name		xargs -n1 rmdev -dl	for 2107 devices
lsdev -C -t 1750* -F name		xargs -n1 rmdev -dl	for 1750 devices
 - b. Verify that the hdisk devices are successfully removed using the following command:

lsdev -C -t 2105* -F name	for 2105 devices
lsdev -C -t 2145* -F name	for 2145 devices
lsdev -C -t 2107* -F name	for 2107 devices
lsdev -C -t 1750* -F name	for 1750 devices
4. If you are upgrading the operating system, follow these procedures. Otherwise, if you are not upgrading the operating system, skip to step 5 on page 33.
 - a. Run **stopsrc -s sddsrv** to stop the sddsrv daemon.
 - b. Uninstall SDD.

- c. Upgrade to the latest version of the host attachment, if required. The following are package names:
 - `ibm2105.rte` for 2105 devices
 - `devices.fcp.disk.ibm.rte` for 2145, 2107, and 1750 devices
 - d. If `rootvg` is on a SAN boot disk, restart the system.
 - e. Make sure no disk group is online except `rootvg`. Migrate the AIX OS level. The system automatically restarts at the end of migration.
 - f. Install SDD for the new AIX OS level.
 - g. Configure SDD vpath devices by running the `cfallvpath` command.
 - h. Continue to step 6.
5. If you are not upgrading the operating system, follow these steps.
 - a. Upgrade to the latest version of Host Attachment, if required. The following are Host Attachment Package names:
 - `ibm2105.rte` for 2105 devices
 - `devices.fcp.disk.ibm.rte` for 2145, 2107, and 1750 devices
 - b. After upgrading Host Attachment,
 - If `rootvg` is on a SAN boot disk, restart the system. Then skip the rest of the steps and follow the procedures in “Upgrading the SDD packages automatically without system restart” on page 24 to upgrade SDD, if required.
 - If `rootvg` is on local SCSI disks and you can restart the system, skip the rest of the steps and restart the system. Then follow the procedures in “Upgrading the SDD packages automatically without system restart” on page 24 to upgrade SDD, if required.
 - If `rootvg` is on local SCSI disks and you cannot restart the system, continue to the next step.
 - c. Upgrade to the latest version of SDD, if required.
 - d. Configure hdisks and SDD vpath devices by running the `cfgmgr` command.
 6. If your new SDD version is earlier than 1.6.0.0, run the `hd2vp` command on all SDD volume groups. Otherwise, skip this step.
 7. Resume all activities related to SDD devices:
 - a. If there was active paging space created with SDD devices, activate the paging space.
 - b. If your host was in an HACMP environment, start the cluster services.
 - c. Vary on all SDD volume groups.
 - d. Mount all file systems.
 - e. Start applications running on SDD volume groups or file systems.

Case 2: In HACMP node-by-node migration with nonconcurrent resource group

If you are upgrading the software on the HACMP standby node during the HACMP node-by-node migration with non-concurrent resource group, follow these procedures.

Note: Currently, migration to AIX 5.2 with NIM is not supported.

On the HACMP standby node, complete the following steps to upgrade the software:

1. Stop all activities related to SDD devices, if there are any.
 - a. Stop applications running on SDD volume groups or file systems.

- b. If there is active paging space created with SDD devices, deactivate the paging space.
- c. Use the **umount** command to unmount all file systems of SDD volume groups.
- d. Vary off all SDD volume groups.

Ensure that all SDD vpath devices are closed by running the **datapath query device** command to verify that all the paths of all SDD vpath devices are at CLOSE state.

2. Remove the SDD vpath devices using the **rmdev -dl dpo -R** command. See “Special considerations when you uninstall SDD in the NIM environment” on page 50.
3. Run **stopsrc -s sddsrv** to stop the sddsrv daemon.
4. Uninstall SDD.
5. If you do not want to upgrade OS, skip to step 10 on page 35. If you want to upgrade OS, go to the step 6.
6. If *rootvg* is on local SCSI disks, follow these procedures. Otherwise, if *rootvg* is on a SAN boot disk, skip to step 7.

- a. Remove hdisk devices using the following command:

```
lsdev -C -t 2105* -F name | xargs -n1 rmdev -dl      for 2105 devices
lsdev -C -t 2145* -F name | xargs -n1 rmdev -dl      for 2145 devices
lsdev -C -t 2107* -F name | xargs -n1 rmdev -dl      for 2107 devices
lsdev -C -t 1750* -F name | xargs -n1 rmdev -dl      for 1750 devices
```

- b. Verify that the hdisk devices are successfully removed using the following command:

```
lsdev -C -t 2105* -F name      for 2105 devices
lsdev -C -t 2145* -F name      for 2145 devices
lsdev -C -t 2107* -F name      for 2107 devices
lsdev -C -t 1750* -F name      for 1750 devices
```

7. Make sure no disk group is online except *rootvg*. Migrate to the desired AIX OS level. Ensure you complete the following operations for the OS migration.
 - If you are using NIM to upgrade to AIX 5.3, make sure NIM SPOT contains AIX Interim Fix APAR IY94507.
 - Change the option to automatically import user volume groups to **no**.

Reboot automatically starts at the end of migration.

8. If *rootvg* is on a local SCSI disk, follow these procedures. Otherwise, if *rootvg* is on a SAN boot disk, skip to Step 9.

- a. Remove all the hdisks of the SDD supported storage devices with the following command.

```
lsdev -C -t 2105* -F name | xargs -n1 rmdev -dl      for 2105 devices
lsdev -C -t 2145* -F name | xargs -n1 rmdev -dl      for 2145 devices
lsdev -C -t 2107* -F name | xargs -n1 rmdev -dl      for 2107 devices
lsdev -C -t 1750* -F name | xargs -n1 rmdev -dl      for 1750 devices
```

- b. Verify that the hdisk devices are successfully removed using the following command:

```
lsdev -C -t 2105* -F name      for 2105 devices
lsdev -C -t 2145* -F name      for 2145 devices
lsdev -C -t 2107* -F name      for 2107 devices
lsdev -C -t 1750* -F name      for 1750 devices
```

9. Upgrade to the latest version of Host Attachment, if required. The following are Host Attachment Package names:
 - *ibm2105.rte* for 2105 devices
 - *devices.fcp.disk.ibm.rte* for 2145, 2107, and 1750 devices

10. If *rootvg* is on a SAN boot disk, restart the system.
11. Install SDD. If you have migrated to a new AIX OS level, make sure you install the SDD for the new AIX OS level.
12. On the HACMP active node, run **varyonvg -bu *volume group name*** on all the SDD non-concurrent volume groups that are shared with the standby node.
13. On the HACMP standby node, complete the following steps:
 - a. Configure hdisks and the SDD vpath devices using one of the following options:
 - Run **cfgmgr -vl fcsX** for each fibre channel adapter and then run **cfallvpath**
 - Run **cfgmgr**
 - b. If your new SDD version is earlier than 1.6.0.0, run **hd2vp** on all SDD volume groups. Otherwise, skip this step.
 - c. Run **importvg -L *volume group name physical volume name*** to update any possible Object Data Manager (ODM) changes on a volume group.
14. On the HACMP active node, run **varyonvg *volume group name*** on all SDD non-concurrent volume groups that are shared with the standby node.

Verifying the currently installed version of SDD for SDD 1.3.3.11 (or earlier)

For SDD packages prior to SDD 1.4.0.0, you can verify your currently installed version of SDD by entering the following command:

```
lslpp -l '*Sdd*'
```

The asterisks (*) in the beginning and end of the Sdd characters are used as wildcard symbols to search for the characters "ibm..." and "...rte".

Alternatively, you can enter one of the following commands:

```
lslpp -l ibmSdd_432.rte
```

```
lslpp -l ibmSdd_433.rte
```

```
lslpp -l ibmSdd_510.rte
```

```
lslpp -l ibmSdd_510nchacmp.rte
```

```
lslpp -l ibmSdd.rte.432
```

...

...

If you successfully installed the package, the output from the `lslpp -l '*Sdd*'` or `lslpp -l ibmSdd_432.rte` command looks like this:

Fileset	Level	State	Description
Path: /usr/lib/objrepos ibmSdd_432.rte	1.3.3.9	COMMITTED	IBM SDD AIX V432 V433 for concurrent HACMP
Path: /etc/objrepos ibmSdd_432.rte	1.3.3.9	COMMITTED	IBM SDD AIX V432 V433 for concurrent HACMP

If you successfully installed the `ibmSdd_433.rte` package, the output from the `lslpp -l ibmSdd_433.rte` command looks like this:

Fileset	Level	State	Description
Path: /usr/lib/objrepos ibmSdd_433.rte	1.3.3.9	COMMITTED	IBM SDD AIX V433 for nonconcurrent HACMP
Path: /etc/objrepos ibmSdd_433.rte	1.3.3.9	COMMITTED	IBM SDD AIX V433 for nonconcurrent HACMP

If you successfully installed the `ibmSdd_510.rte` package, the output from the `lslpp -l ibmSdd_510.rte` command looks like this:

Fileset	Level	State	Description
Path: /usr/lib/objrepos ibmSdd_510.rte	1.3.3.9	COMMITTED	IBM SDD AIX V510 for concurrent HACMP
Path: /etc/objrepos ibmSdd_510.rte	1.3.3.9	COMMITTED	IBM SDD AIX V510 for concurrent HACMP

If you successfully installed the `ibmSdd_510nchacmp.rte` package, the output from the `lslpp -l ibmSdd_510nchacmp.rte` command looks like this:

Fileset	Level	State	Description
Path: /usr/lib/objrepos ibmSdd_510nchacmp.rte	1.3.3.11	COMMITTED	IBM SDD AIX V510 for nonconcurrent HACMP
Path: /etc/objrepos ibmSdd_510nchacmp.rte	1.3.3.11	COMMITTED	IBM SDD AIX V510 for nonconcurrent HACMP

Verifying the currently installed version of SDD for SDD 1.4.0.0 (or later)

For SDD 1.4.0.0 (and later), you can verify your currently installed version of SDD by entering the following command:

```
lslpp -l 'devices.sdd.*'
```

Alternatively, you can enter one of the following commands:

```
lslpp -l devices.sdd.43.rte
lslpp -l devices.sdd.51.rte
lslpp -l devices.sdd.52.rte
lslpp -l devices.sdd.53.rte
lslpp -l devices.sdd.61.rte
```

If you successfully installed the `devices.sdd.43.rte` package, the output from the `lslpp -l 'devices.sdd.*'` command or `lslpp -l devices.sdd.43.rte` command looks like this:

Fileset	Level	State	Description
Path: /usr/lib/objrepos devices.sdd.43.rte	1.4.0.0	COMMITTED	IBM Subsystem Device Driver for AIX V433
Path: /etc/objrepos devices.sdd.43.rte	1.4.0.0	COMMITTED	IBM Subsystem Device Driver for AIX V433

If you successfully installed the `devices.sdd.51.rte` package, the output from the `lslpp -l devices.sdd.51.rte` command looks like this:

Fileset	Level	State	Description
Path: /usr/lib/objrepos devices.sdd.51.rte	1.4.0.0	COMMITTED	IBM Subsystem Device Driver for AIX V51
Path: /etc/objrepos devices.sdd.51.rte	1.4.0.0	COMMITTED	IBM Subsystem Device Driver for AIX V51

If you successfully installed the `devices.sdd.52.rte` package, the output from the `lslpp -l devices.sdd.52.rte` command looks like this:

Fileset	Level	State	Description
Path: /usr/lib/objrepos devices.sdd.52.rte	1.4.0.0	COMMITTED	IBM Subsystem Device Driver for AIX V52
Path: /etc/objrepos devices.sdd.52.rte	1.4.0.0	COMMITTED	IBM Subsystem Device Driver for AIX V52

If you successfully installed the `devices.sdd.53.rte` package, the output from the `lslpp -l devices.sdd.53.rte` command looks like this:

Fileset	Level	State	Description
Path: /usr/lib/objrepos devices.sdd.53.rte	1.6.0.0	COMMITTED	IBM Subsystem Device Driver for AIX V53
Path: /etc/objrepos devices.sdd.53.rte	1.6.0.0	COMMITTED	IBM Subsystem Device Driver for AIX V53

If you successfully installed the `devices.sdd.61.rte` package, the output from the `lslpp -l devices.sdd.61.rte` command looks like this:

Fileset	Level	State	Description
Path: /usr/lib/objrepos devices.sdd.61.rte	1.7.0.0	COMMITTED	IBM Subsystem Device Driver for AIX V61
Path: /etc/objrepos devices.sdd.61.rte	1.7.0.0	COMMITTED	IBM Subsystem Device Driver for AIX V61

Preparing to configure SDD

Before you configure SDD, ensure that:

- The supported storage device is operational.
- The `devices.sdd.mm.rte` software is installed on the AIX host system, where *mm* identifies the installation package.
- The supported storage device hdisks are configured correctly on the AIX host system.

Configure the supported storage devices before you configure SDD. If you configure multiple paths to a supported storage device, ensure that all paths (hdisks) are in Available state. Otherwise, some SDD vpath devices will lose multipath capability.

Complete the following steps:

1. Enter the `lsdev -C -t xxxx` command to check the supported storage hdisk device configuration, where *xxxx* is the supported storage device type.
You can pipe the output of the `lsdev` command to `grep` for a certain type of device. For example, use one of the following commands:
 - `lsdev -C -t 2105` to check the ESS device configuration
 - `lsdev -C -t 2107` to check the DS8000 device configuration
 - `lsdev -C -t 1750` to check the DS6000 device configuration
 - `lsdev -C -t 2145` to check the SAN Volume Controller device configuration
2. If you have already created some active volume groups with SDD supported storage devices, vary off (deactivate) all these active volume groups by using the `varyoffvg` (LVM) command. If there are file systems of these volume groups that are mounted, you must also unmount all file systems in order to configure SDD vpath devices correctly.

Maximum number of LUNs

For different AIX OS levels, SDD has set different limits on the maximum number of LUNs that can be configured. These limits exist because AIX has resource limitations on the total number of devices that a system can support. In a multipath configuration environment, AIX creates one hdisk device for each path to a physical disk. Increasing the number of paths that are configured to a physical disk increases the number of AIX system hdisk devices that are created and are consuming system resources. This might leave fewer resources for SDD vpath devices to be configured. On the other hand, more SDD vpath devices can be configured if the number of paths to each disk is reduced.

For AIX versions 4.3 and 5.1, AIX has a published limit of 10 000 devices per system. Based on this limitation, SDD limits the total maximum number of SDD vpath devices that can be configured to 600. This number is shared by all SDD-supported storage devices.

For AIX version 5.2 or later, the resource of the AIX operating system is increased. SDD has increased the SDD vpath device limit accordingly. Beginning with SDD 1.6.0.7, SDD supports a combined maximum of 1200 supported storage devices on AIX version 5.2 or later.

Table 7 on page 39 provides a summary of the maximum number of LUNs and the maximum number of paths allowed when running on a host systems with different operating system levels.

Table 7. Maximum LUNs allowed for different AIX OS levels

OS level	SDD supported storage devices
AIX 4.3 *	600 LUNs (maximum 32 paths)
AIX 5.1	600 LUNs (maximum 32 paths)
AIX 5.2	1200 LUNs (maximum 32 paths; see Table 8 for recommended maximum number of paths.)
AIX 5.3	1200 LUNs (maximum 32 paths; see Table 8 for recommended maximum number of paths.)
AIX 6.1	1200 LUNs (maximum 32 paths; see Table 8 for recommended maximum number of paths.)

* **Note:** AIX 4.3 is only supported for ESS and virtualization products.

You can have a maximum of 32 paths per SDD vpath device regardless of the number of LUNs that are configured. This means that you can only have a maximum of 32 host adapter ports for your SDD vpath devices. However, configuring more paths than is needed for failover protection might consume too many system resources and degrade system performance. Use the minimum number of paths that are necessary to achieve sufficient redundancy in the SAN environment. The recommended number of paths is 2 - 4.

To avoid exceeding the maximum number of paths per SDD vpath device on AIX 5.2 or later, follow the recommendations in Table 8.

Table 8. Recommended maximum paths supported for different number of LUNs on AIX 5.2 or later

Number of LUNs	Maximum paths per vpath
1- 600 vpath LUN	16
601 - 900 vpath LUN	8
901 - 1200 vpath LUN*	4

Note: * In order to configure 1200 LUNs, APAR IY49825 is required.

If you have more than 1200 vpaths already configured on your AIX host (for example, if you have 800 ESS LUNs and 512 SAN Volume Controller LUNs configured as SDD vpath devices on one AIX host), SDD migration to SDD 1.6.0.7 or later will fail because SDD does not support more than 1200 LUNs on one AIX host. If you have this configuration, contact IBM Customer Support at 1-800-IBM-SERV.

ODM attributes for controlling the maximum number of LUNs in SDD version 1.6.0.7 or later on AIX 5.2 and later

SDD for AIX 5.2 and later has consolidated the ODM attributes for controlling the maximum number of LUNs for all supported storage devices. The SDD_maxlun ODM attribute is now used to replace the following ODM attributes:

- 2105_max_luns
- 2145_max_luns
- 2062_max_luns
- Enterpr_maxlun
- Virtual_maxlun

See Table 7 on page 39 for information about the total number of LUNs that you can configure.

The new SDD ODM attribute, `SDD_maxlun`, defines the maximum number of storage LUNs that SDD can support on a host. This attribute has a default value as well as a maximum value of 1200. This value is not user-changeable.

To display the value of the `SDD_maxlun` attribute, use the `lsattr -El dpo` command:

```
> lsattr -El dpo
SDD_maxlun 1200          Maximum LUNS allowed for SDD          False
persistent_resv yes    Subsystem Supports Persistent Reserve Command False
```

Preparing your system to configure more than 600 supported storage devices or to handle a large amount of I/O after queue depth is disabled

If you plan to configure more than 600 supported storage devices by configuring multiple types of supported storage systems and the total number of LUNs will exceed 600, or if you plan to disable queue depth to remove the limit on the amount of I/O that SDD vpath devices can send, you must first determine whether the system has sufficient resources for large device configuration or heavy I/O operations.

There are also some system configurations that must be changed to avoid system bottleneck.

To avoid system-performance degradation, tune the following ODM attributes for your AIX fibre-channel adapters before you configure more than 600 supported storage devices or disable queue depth:

- `lg_term_dma`
- `num_cmd_elems`
- `max_xfer_size`
- `fc_err_recov`

If you change these attributes, you need to reconfigure the fibre-channel adapter and all its child devices. Because this is a disruptive procedure, change these attributes before assigning or configuring supported storage devices on a host system.

lg_term_dma

This AIX fibre-channel adapter attribute controls the DMA memory resource that an adapter driver can use. The default value of `lg_term_dma` is 0x200000, and the maximum value is 0x8000000. A recommended change is to increase the value of `lg_term_dma` to 0x400000. If you still experience poor I/O performance after changing the value to 0x400000, you can increase the value of this attribute again. If you have a dual-port fibre-channel adapter, the maximum value of the `lg_term_dma` attribute is divided between the two adapter ports. Therefore, never increase `lg_term_dma` to the maximum value for a dual-port fibre-channel adapter, because this will cause the configuration of the second adapter port to fail.

num_cmd_elems

This AIX fibre-channel adapter attribute controls the maximum number of commands to be queued to the adapter. The default value is 200, and the maximum value is:

LP9000 adapters
2048

LP10000 adapters

2048

LP7000 adapters

1024

When a large number of supported storage devices are configured, you can increase this attribute to improve performance.

max_xfer_size

This AIX fibre-channel adapter attribute controls the maximum transfer size of the fibre-channel adapter. Its default value is 100000 and the maximum value is 1000000. You can increase this attribute to improve performance. Different storages might need different maximum transfer size to utilize the performance.

Note: You can change this attribute only with AIX 5.2.0 or later.

fc_err_recov

Beginning with AIX 5.1 and AIX52 TL02, the `fc_err_recov` attribute enables fast failover during error recovery. Enabling this attribute can reduce the amount of time that the AIX disk driver takes to fail I/O in certain conditions, and therefore, reduce the overall error recovery time. The default value for `fc_err_recov` is *delayed_fail*. To enable fibre-channel adapter fast failover, change the value to *fast_fail*.

Note:

1. For AIX 5.1, apply APAR IY48725 (Fast I/O Failure for Fibre Channel Devices) to add the fast failover feature.
2. Fast failover is not supported on AIX 4.3.3 or earlier.

Using the **-P** option with the **chdev** command causes the changes to not take effect until the system is restarted. Use the **-P** option with the **chdev** command if your system falls under any one of the following conditions:

- If you have boot devices on the adapter
- If you have a large number of devices configured and prefer to restart the system
- If you plan to reboot the system later

Use the following procedure if you can reboot the system and allow the new attribute values to take effect after the reboot:

1. Issue `lsattr -El fcsN` to check the current value of `lg_term_dma`, `num_cmd_elems`, and `max_xfer_size`.
2. Issue `lsattr -El fscsiN` to check the current value of `fc_err_recov`.
3. Issue `chdev -l fcsN -P -a lg_term_dma=0x400000` to increase the DMA value.
4. Issue `chdev -l fcsN -P -a num_cmd_elems=1024` to increase the maximum commands value.
5. Issue `chdev -l fcsN -P -a max_xfer_size=20000` to increase the maximum transfer size.
6. Issue the `chdev -l fscsiX -P -a fc_err_recov=fast_fail` command to enable fast failover.
7. Assign new LUNs to the AIX host, if needed.
8. Reboot the system now or later.

Use the following procedure if you cannot reboot the system but want the new attributes to take effect immediately:

1. Issue **lsattr -El fcsN** to check the current value of `lg_term_dma`, `num_cmd_elems`, and `max_xfer_size`.
2. Issue **lsattr -El fscsiN** to check the current value of `fc_err_recov`.
3. Use the **rmdev -dl dpo -R** command to remove SDD vpath devices, if they are already configured on your system.
4. Put all existing fibre-channel adapters and their child devices to the Defined state by issuing **rmdev -l fcsN -R**.
5. Issue **chdev -l fcsN -a lg_term_dma=0x400000** to increase the DMA value.
6. Issue **chdev -l fcsN -a num_cmd_elems=1024** to increase the maximum commands value.
7. Issue **chdev -l fcsN -a max_xfer_size=100000** to increase the maximum transfer size.
8. Issue **chdev -l fscsiX -a fc_err_recov=fast_fail** to enable fast failover.
9. Assign new LUNs to the AIX host, if needed.
10. Configure the fibre-channel adapters, the child devices and hdisks using **cfgmgr -l fcsN**.
11. Configure SDD vpath devices with the **cfallvpath** command if they are removed in step 3.

When you have a large number of LUNs, many special device files are created in the `/dev` directory. Issuing the **ls** command with a wildcard (*) in this directory might fail. If issuing the command fails in this situation, change the `ncargs` attribute of `sys0`. The `ncargs` attribute controls the ARG/ENV list size in 4-KB byte blocks. The default value for this attribute is 6 (24 KB) and the maximum value for this attribute is 128 (512 KB). Increase the value of this to 30. If you still experience failures after changing the value to 30, increase this value to a larger number. Changing the `ncargs` attribute is dynamic. Use the following command to change the `ncargs` attribute to 30:

```
chdev -l sys0 -a ncargs=30
```

Filesystem space:

If you are increasing the maximum number of LUNs, after changing the ODM attributes, use the following steps to determine whether there is sufficient space in the root file system after hdisks are configured:

1. Issue **cfgmgr -l [scsiN/fcsN]** for each relevant SCSI or FCP adapter.
2. Issue **df** to ensure that root file system (that is, '/') size is large enough to hold the device special files. For example:

```
Filesystem 512-blocks  Free %Used  Iused  %Iused  Mounted on
/dev/hd4    196608    29008  86%    15524  32%     /
```

The minimum required size is 8 MB. If there is insufficient space, run the **chfs** command to increase the size of the root file system.

Controlling I/O flow to SDD devices with the SDD `qdepth_enable` attribute

Starting with SDD 1.5.0.0, the SDD attribute, `qdepth_enable`, allows you to control I/O flow to SDD vpath devices. `qdepth_enable` was a `dpo` attribute before SDD 1.6.1.2 and it controls the queue depth logic on all the SDD vpath devices.

Beginning with SDD 1.6.1.2, `qdepth_enable` is changed to a vpath attribute and it allows you to set different queue depth logic per SDD vpath device.

By default, SDD uses the device `queue_depth` setting to control the I/O flow to SDD vpath device and paths. With certain database applications, such as an application running with a DB2® database, IBM Lotus® Notes®, or IBM Informix® database, the software might generate many threads, which can send heavy I/O to a relatively small number of devices. Enabling queue depth logic to control I/O flow can cause performance degradation, or even a system hang. To remove the limit on the amount of I/O sent to vpath devices, use the `qdepth_enable` attribute to disable this queue depth logic on I/O flow control.

By default, the queue depth logic to control the amount of I/O being sent to the vpath devices is enabled in the SDD driver.

To determine if queue depth logic is enabled for a particular SDD vpath device, run the following command:

```
# lsattr -El vpath0
active_hdisk hdisk66/13AB2ZA1020/fscsi3      Active hdisk      False
active_hdisk hdisk2/13AB2ZA1020/fscsi2      Active hdisk      False
active_hdisk hdisk34/13AB2ZA1020/fscsi2     Active hdisk      False
active_hdisk hdisk98/13AB2ZA1020/fscsi3     Active hdisk      False
policy       df                               Scheduling Policy True
pvid         0005f9fdcda4417d000000000000000 Physical volume identifier False
qdepth_enable yes                          Queue Depth Control True
reserve_policy PR_exclusive                 Reserve Policy     True
serial_number 13AB2ZA1020                   LUN serial number False
unique_id     yes                           Device Unique Identification False
```

For SDD 1.5.1.0 or later, you can change the `qdepth_enable` attribute dynamically. The `datapath set qdepth` command offers a new option to dynamically enable or disable the queue depth logic. For example, if you enter `datapath set device 0 2 qdepth disable` command, the following output is displayed when the queue depth logic is currently enabled on these SDD vpath devices:

```
Success: set qdepth_enable to no for vpath0
Success: set qdepth_enable to no for vpath1
Success: set qdepth_enable to no for vpath2
```

The `qdepth_enable` ODM attribute of these SDD vpath devices will be updated. For example, the following output is displayed when `lsattr -El vpath0` is entered.

```
# lsattr -El vpath0
active_hdisk hdisk66/13AB2ZA1020/fscsi3      Active hdisk      False
active_hdisk hdisk2/13AB2ZA1020/fscsi2      Active hdisk      False
active_hdisk hdisk34/13AB2ZA1020/fscsi2     Active hdisk      False
active_hdisk hdisk98/13AB2ZA1020/fscsi3     Active hdisk      False
policy       df                               Scheduling Policy True
pvid         0005f9fdcda4417d0000000000000000 Physical volume identifier False
qdepth_enable no                          Queue Depth Control True
reserve_policy PR_exclusive                 Reserve Policy     True
serial_number 13AB2ZA1020                   LUN serial number False
unique_id     yes                           Device Unique Identification False
```

See “Preparing your system to configure more than 600 supported storage devices or to handle a large amount of I/O after queue depth is disabled” on page 40 to determine whether the system has sufficient resources for disabling queue depth logic.

Controlling reserve policy of SDD devices with the SDD reserve_policy attribute

Starting with SDD 1.7.1.0, a new **reserve_policy** attribute allows you to control the default reserve policy of each SDD vpath device.

By default, the **reserve_policy** attribute is set to **PR_exclusive**. This means that, unless an application opens the SDD vpath device with the **no reserve device open** option, SDD always makes a persistent reservation on the device.

In a shared environment where an application does not implement the **no reserve device open** option, such as the Dual Virtual I/O Servers environment, you must set the **reserve_policy** attribute to **no_reserve**. If you set the **reserve_policy** attribute to **no_reserve**, regardless of the device open option, SDD does not make a persistent reservation on the device.

To display the current **reserve_policy**, run the following command:

```
# lsattr -El vpath0
active_hdisk hdisk331/140FCA30/fscsi1 Active hdisk False
active_hdisk hdisk2/140FCA30/fscsi0 Active hdisk False
active_hdisk hdisk53/140FCA30/fscsi0 Active hdisk False
active_hdisk hdisk197/140FCA30/fscsi0 Active hdisk False
active_hdisk hdisk280/140FCA30/fscsi1 Active hdisk False
active_hdisk hdisk475/140FCA30/fscsi1 Active hdisk False
policy df Scheduling Policy True
pvid 00082dbaldae728c0000000000000000 Physical volume identifier False
qdepth_enable yes Queue Depth Control True
reserve_policy PR_exclusive Reserve Policy True
serial_number 140FCA30 LUN serial number False
unique_id yes Device Unique Identification False
```

To change the vpath attribute, **reserve_policy**, from default value “**PR_exclusive**” to “**no_reserve**”, enter the following command:

```
# chdev -l vpath0 -a reserve_policy=no_reserve
vpath0 changed
```

The **chdev** command requires that the device is closed because it reconfigures the SDD vpath device.

Configuring SDD

Complete the following steps to configure SDD using SMIT:

Note: The list items on the SMIT panel might be worded differently from one AIX version to another.

1. Enter **smitty device** from your desktop window. The Devices menu is displayed.
2. Select **Data Path Device** and press **Enter**. The Data Path Device panel is displayed.
3. Select **Define and Configure All Data Path Devices** and press **Enter**. The configuration process begins.
4. Check the SDD configuration state. See “Displaying the supported storage device SDD vpath device configuration” on page 73.
5. Use the **varyonvg** command to vary on all deactivated supported storage device volume groups.
6. Mount the file systems for all volume groups that were previously unmounted.

Unconfiguring SDD

1. Before you unconfigure SDD devices, ensure that:
 - All I/O activities on the devices that you need to unconfigure are stopped.
 - All file systems belonging to the SDD volume groups are unmounted and all volume groups are varied off.
 - A paging space created with SDD devices is deactivated.
2. Run the `vp2hd volume_group_name` conversion script to convert the volume group from SDD devices (vpathN) to supported storage devices (hdisks).

Note: Because SDD implements persistent reserve command set, you must remove the SDD vpath device before removing the SDD vpath device's underlying hdisk devices.

You can use SMIT to unconfigure the SDD devices in two ways. You can either unconfigure *without deleting* the device information from the Object Database Manager (ODM) database, or you can unconfigure and *delete* device information from the ODM database:

- If you unconfigure without deleting the device information, the device remains in the *Defined* state. You can use either SMIT or the `mkdev -l vpathN` command to return the device to the *Available* state.
- If you unconfigure and delete the device information from the ODM database, that device is removed from the system. To reconfigure it, follow the procedure described in “Configuring SDD” on page 44.

Complete the following steps to delete device information from the ODM and to unconfigure SDD devices:

1. Enter **smitty device** from your desktop window. The Devices menu is displayed.
2. Select **Devices** and press **Enter**.
3. Select **Data Path Device** and press **Enter**. The Data Path Device panel is displayed.
4. Select **Remove a Data Path Device** and press **Enter**. A list of all SDD devices and their states (either *Defined* or *Available*) is displayed.
5. Select the device that you want to unconfigure. Select whether you want to delete the device information from the ODM database.
6. Press **Enter**. The device is unconfigured to the state that you selected.
7. To unconfigure more SDD devices, you have to repeat steps 4 - 6 for each SDD device.

The fast-path command to unconfigure all SDD devices and change the device state from *Available* to *Defined* is: `rmdev -l dpo -R`. The fast-path command to unconfigure and remove all SDD devices from your system is: `rmdev -dl dpo -R`.

Verifying the SDD configuration

To check the SDD configuration, you can use either the SMIT Display Device Configuration panel or the `lsvpcfg` console command.

Complete the following steps to verify the SDD configuration on an AIX host system:

1. Enter **smitty device** from your desktop window. The Devices menu is displayed.

2. Select **Data Path Device** and press **Enter**. The Data Path Device panel is displayed.
3. Select **Display Data Path Device Configuration** and press **Enter**.
4. Select all devices for Select Query Option, leave the Device Name/ Device Model field blank and press **Enter**. The state (either Defined or Available) of all SDD vpath devices and the paths to each device is displayed.

If any device is listed as *Defined*, the configuration was not successful. Check the configuration procedure again. See “Configuring SDD” on page 44 for the procedure.

If you want to use the command-line interface to verify the configuration, enter **lsvpcfg**.

You should see an output similar to this:

```
vpath0 (Avail pv vpathvg) 018FA067 = hdisk1 (Avail )
vpath1 (Avail ) 019FA067 = hdisk2 (Avail )
vpath2 (Avail ) 01AFA067 = hdisk3 (Avail )
vpath3 (Avail ) 01BFA067 = hdisk4 (Avail ) hdisk27 (Avail )
vpath4 (Avail ) 01CFA067 = hdisk5 (Avail ) hdisk28 (Avail )
vpath5 (Avail ) 01DFA067 = hdisk6 (Avail ) hdisk29 (Avail )
vpath6 (Avail ) 01EFA067 = hdisk7 (Avail ) hdisk30 (Avail )
vpath7 (Avail ) 01FFA067 = hdisk8 (Avail ) hdisk31 (Avail )
vpath8 (Avail ) 020FA067 = hdisk9 (Avail ) hdisk32 (Avail )
vpath9 (Avail pv vpathvg) 02BFA067 = hdisk20 (Avail ) hdisk44 (Avail )
vpath10 (Avail pv vpathvg) 02CFA067 = hdisk21 (Avail ) hdisk45 (Avail )
vpath11 (Avail pv vpathvg) 02DFA067 = hdisk22 (Avail ) hdisk46 (Avail )
vpath12 (Avail pv vpathvg) 02EFA067 = hdisk23 (Avail ) hdisk47 (Avail )
vpath13 (Avail pv vpathvg) 02FFA067 = hdisk24 (Avail ) hdisk48 (Avail )
```

The output shows:

- The name of each SDD vpath device (for example, vpath13)
- The *Defined* or *Available* state of an SDD vpath device
- Whether the SDD vpath device is defined to AIX as a physical volume (indicated by the pv flag)
- The name of the volume group the device belongs to (for example, vpathvg)
- The unit serial number of the disk storage system LUN (for example, 02FFA067) or the unit serial number of the virtualization product LUN (for example, 60056768018A0210B000000000000006B)
- The names of the AIX disk devices making up the SDD vpath device and their configuration and physical volume state

Dynamically adding paths to SDD vpath devices

This section shows you how to add paths to SDD vpath devices with the **addpaths** command.

The **addpaths** command allows you to dynamically add more paths to SDD vpath devices when they are in the *Available* state. It also allows you to add paths to SDD vpath devices that belong to active volume groups.

If an SDD vpath device is opened, the paths that are added are automatically in the Open state. With SDD levels earlier than 1.5.1.0, there is an exception when you enter the **addpaths** command to add a new path to an SDD vpath device that has only one configured path. In this case, the new path is not automatically in the Open state, and you must change it to the Open state by closing and reopening the

SDD vpath device. This exception is removed for SDD 1.5.1.0 and later. That is, in SDD 1.5.1.0 and later, the new path automatically opens after you add it to an opened SDD vpath device.

You can determine how many more hdisks can be added to the existing SDD vpath devices by using the command **datapath query device** to find out how many hdisks are already configured for each vpath. Ensure that the number of hdisks from the existing configuration is below the recommended maximum paths allowed according to Table 8 on page 39.

Note:

1. If you configure more than 32 hdisks, running **addpaths** will not add any paths to vpaths.
2. Before you run **addpaths**, you must run **cfgmgr** to configure the new hdisks devices. Running **cfgmgr** alone does not add new paths to SDD vpath devices.

Dynamically removing or replacing PCI adapters or paths

SDD 1.5.1.0 (or later) supports AIX Hot Plug available on 5L or later. You can dynamically replace an adapter in a hot-plug slot. You can use the AIX **lsslot** command to display dynamically reconfigurable slots, such as hot-plug slots, and their characteristics. You can also remove a particular path of an SDD vpath device. Replacing an adapter or removing paths does not interrupt current I/O and SDD can be dynamically reconfigured without shutting down or powering off the system. Three possible scenarios using this feature in the SDD environment are:

- “Dynamically removing a PCI adapter from SDD configuration”
- “Dynamically replacing a PCI adapter in an SDD configuration”
- “Dynamically removing a path of an SDD vpath device” on page 48

Dynamically removing a PCI adapter from SDD configuration

To permanently remove a PCI adapter and its child devices from an SDD configuration, use the **datapath remove adapter *n*** command, where *n* is the adapter number.

Dynamically replacing a PCI adapter in an SDD configuration

To dynamically replace a PCI adapter in an SDD configuration, use the **datapath remove adapter *n*** command, where *n* is the adapter number. This command removes the adapter and associated paths from the SDD configuration.

After you physically replace and configure a new adapter, the adapter and its associated paths can be added to SDD with the **addpaths** command. See “datapath remove adapter” on page 386 for more information about the **datapath remove adapter *n*** command.

Complete the following steps to dynamically replace a PCI adapter in the SDD configuration:

1. Enter **datapath query adapter** to identify the adapter to be replaced.

```
-----+-----
Active Adapters :4
-----+-----
Adpt#   Adapter Name  State   Mode   Select   Errors  Paths  Active
-----+-----
0       fscsi0         NORMAL ACTIVE  62051    415     10     10
1       fscsi1         NORMAL ACTIVE  65386    3       10     10
2       fscsi2         NORMAL ACTIVE  75697    27      10     10
3       fscsi3         NORMAL ACTIVE  4788     35     10     10
-----+-----
```

- Enter **datapath remove adapter *n***, where *n* is the adapter number to be removed. For example, to remove adapter 0, enter **datapath remove adapter 0**.

```

+-----+
|Success: remove adapter 0
|
|Active Adapters :3
|
|Adpt#   Adapter Name  State   Mode   Select   Errors  Paths   Active
|  1     fscsi1     NORMAL  ACTIVE  65916    3       10      10
|  2     fscsi2     NORMAL  ACTIVE  76197    28      10      10
|  3     fscsi3     NORMAL  ACTIVE  4997     39      10      10
+-----+

```

Note that Adpt# 0 fscsi0 is removed and the Select counts are increased on other three adapters, indicating that I/O is still running.

- Enter **rmdev -dl fcs0 -R** to remove fcs0, a parent of fscsi0, and all of its child devices from the system. Issuing **lsdev -Cc disk** should not show any devices associated with fscsi0.
- Enter **drslot -R -c pci -s P1-I8** where P1-I8 is the slot location found by issuing **lscfg -vl fcs0**. This command prepares a hot-plug slot for systems with AIX 5L™ or later.
- Follow the instruction given by **drslot** to physically remove the adapter and install a new one.
- Update the World Wide Name (WWN) of the new adapter at the device end and in the fabric. For example, for ESS devices, go to the ESS Specialist to update the WWN of the new adapter. The zone information of fabric switches must be updated with the new WWN as well.
- Enter **cfgmgr** or **cfgmgr -vl pci(*n*)**, where *n* is the adapter number, to configure the new adapter and its child devices. Use the **lsdev -Cc disk** and **lsdev -Cc adapter** commands to ensure that all devices are successfully configured to *Available* state.
- Enter the **addpaths** command to configure the newly installed adapter and its child devices to SDD. The newly added paths are automatically opened if vpath is open.

```

+-----+
|Active Adapters :4
|
|Adpt#   Adapter Name  State   Mode   Select   Errors  Paths   Active
|  0     fscsi0     NORMAL  ACTIVE    11        0       10      10
|  1     fscsi1     NORMAL  ACTIVE  196667    6       10      10
|  2     fscsi2     NORMAL  ACTIVE  208697   36       10      10
|  3     fscsi3     NORMAL  ACTIVE   95188    47       10      10
+-----+

```

Dynamically removing a path of an SDD vpath device

To dynamically remove a particular path from an SDD vpath device, use the **datapath remove device *m* path *n*** command. This command permanently removes the logical path from the SDD device. See “datapath remove device path” on page 387 for more information about the **datapath remove device *m* path *n*** command.

Complete the following steps to remove a path of an SDD vpath device:

- Enter **datapath query device** to identify which path of which device is to be removed.

```

+-----+
|DEV#:  0  DEVICE NAME: vpath0  TYPE: 2105E20  POLICY: Optimized
|SERIAL: 20112028
|
|=====
|Path#   Adapter/Hard Disk  State   Mode   Select   Errors
|  0     fscsi1/hdisk18    OPEN    NORMAL  557       0
|  1     fscsi1/hdisk26    OPEN    NORMAL  568       30
|  2     fscsi0/hdisk34    OPEN    NORMAL  566       0
|  3     fscsi0/hdisk42    OPEN    NORMAL  545       0
+-----+

```

2. Enter **datapath remove device *m* path *n***, where *m* is the device number and *n* is the path number of that device. For example, enter **datapath remove device 0 path 1** to remove Path#1 from DEV#0.

```

-----
Success: device 0 path 1 removed
-----
DEV#: 0  DEVICE NAME: vpath0  TYPE: 2105E20  POLICY: Optimized
SERIAL: 20112028
-----
Path#      Adapter/Hard Disk  State   Mode    Select   Errors
-----
0          fscsi1/hdisk18    OPEN   NORMAL  567      0
1          fscsi0/hdisk34    OPEN   NORMAL  596      0
2          fscsi0/hdisk42    OPEN   NORMAL  589      0
-----

```

Note that fscsi1/hdisk26 is removed and Path# 1 is now fscsi0/hdisk34.

3. To reclaim the removed path, see “Dynamically adding paths to SDD vpath devices” on page 46.

```

-----
DEV#: 0  DEVICE NAME: vpath0  TYPE: 2105E20  POLICY: Optimized
SERIAL: 20112028
-----
Path#      Adapter/Hard Disk  State   Mode    Select   Errors
-----
0          fscsi1/hdisk18    OPEN   NORMAL  588      0
1          fscsi0/hdisk34    OPEN   NORMAL  656      0
2          fscsi0/hdisk42    OPEN   NORMAL  599      0
3          fscsi1/hdisk26    OPEN   NORMAL   9       0
-----

```

Note that fscsi1/hdisk26 is added with Path# 3.

Removing SDD from an AIX host system

The SDD server (sddsrv) is an integrated component of SDD 1.3.2.9 (or later). The SDD server daemon is automatically started after SDD is installed. You must stop the SDD server if it is running in the background *before removing SDD*. Go to “Verifying if the SDD server has started” on page 67 and “Stopping the SDD server” on page 68 for more instructions. See Chapter 9, “Using the SDDDSM or the SDDPCM daemon,” on page 365 for more details about the SDD server daemon.

Follow these procedures to completely remove SDD package from your AIX host system.

1. Stop all activities that are related to SDD devices, if there are any.
 - a. Stop applications that are running on SDD volume groups or file systems.
 - b. If there is active paging space that is created with SDD devices, deactivate the paging space.
 - c. Issue the **umount** command to unmount all file systems of SDD volume groups.
 - d. Vary off all SDD volume groups.
2. Remove the SDD vpath devices using the **rmdev -dl dpo -R** command. See “Special considerations when you uninstall SDD in the NIM environment” on page 50.
3. SDD has a server daemon running in the background. Therefore, before uninstalling SDD, issue the **stopsrc -s sddsrv** command to stop the sddsrv daemon.
4. Complete the following steps to uninstall the SDD package.
 - a. Enter **smitty deinstall** from your desktop window to go directly to the Remove Installed Software panel.

- b. Enter one of the following installation package names in the **SOFTWARE name** field:

devices.sdd.43.rte

devices.sdd.51.rte

devices.sdd.52.rte

devices.sdd.53.rte

devices.sdd.61.rte

Then press **Enter**.

Note: See “Verifying the currently installed version of SDD for SDD 1.3.3.11 (or earlier)” on page 35 or “Verifying the currently installed version of SDD for SDD 1.4.0.0 (or later)” on page 36 to verify your currently installed installation package or version of SDD. You can also press **F4** in the Software name field to list the currently installed installation package and do a search (/) on SDD.

- c. Press the **Tab** key in the **PREVIEW Only?** field to toggle between Yes and No. Select **No** to remove the software package from your AIX host system.

Note: If you select **Yes**, the deinstall process will do a pre-check and let you preview the results without removing the software. If the state for any SDD device is either *Available* or *Defined*, the process fails.

- d. Select **No** for the remaining fields on this panel.
e. Press **Enter**. SMIT responds with the following message:

```
ARE YOU SURE??  
Continuing may delete information you may want to keep.  
This is your last chance to stop before continuing.
```

- f. Press **Enter** to begin the removal process. This might take a few minutes.
g. When the process is complete, the SDD software package is removed from your system.

Special considerations when you uninstall SDD in the NIM environment

When you uninstall SDD in the Network Installation Manager (NIM) environment, SDD does not require that you remove all the SDD vpath devices or stop sddsrv daemon before you perform the SDD uninstallation. For example, if you are using NIM to upgrade AIX OS from 5.3 to 6.1, you can add SDD for AIX 5.3 package uninstallation as one of the procedures in your NIM pre-migration script. You do not need to issue the **rmdev** command in your NIM pre-migration script before you uninstall SDD.

Removing SDD Host Attachment from an AIX host system

Follow these procedures to completely remove the SDD Host Attachment package from your AIX host system:

1. Remove hdisk devices using the following command:

```
lsdev -C -t 2105* -F name | xargs -n1 rmdev -dl for 2105 devices  
lsdev -C -t 2145* -F name | xargs -n1 rmdev -dl for 2145 devices  
lsdev -C -t 2107* -F name | xargs -n1 rmdev -dl for 2107 devices  
lsdev -C -t 1750* -F name | xargs -n1 rmdev -dl for 1750 devices
```

2. Verify that the hdisk devices are successfully removed using the following command:

```
lsdev -C -t 2105* -F name for 2105 devices
lsdev -C -t 2145* -F name for 2145 devices
lsdev -C -t 2107* -F name for 2107 devices
lsdev -C -t 1750* -F name for 1750 devices
```

3. Enter **smitty deinstall** from your desktop window to go directly to the Remove Installed Software panel.
4. Enter the following installation package names in the SOFTWARE name field:
 - a. `ibm2105.rte`
 - b. `devices.fcp.disk.ibm.rte`

Note: You can also press F4 in the Software name field to list the currently installed installation package and search (/) on **ibm2105** and **devices.fcp.disk.ibm**.

5. Press the **Tab** key in the **PREVIEW Only?** field to toggle between Yes and No. Select No to remove the software package from your AIX host system.

Note: If you select **Yes**, the deinstall process does a pre-check and lets you preview the results without removing the software. If the state for any SDD device is either Available or Defined, the process fails.

6. Select **No** for the remaining fields on this panel.
7. Press **Enter**. SMIT responds with the following message:

ARE YOU SURE?
Continuing may delete information you may want to keep.
This is your last chance to stop before continuing.

8. Press **Enter** to begin the removal process. This might take a few minutes.
9. When the process is complete, the SDD software package is removed from your system.

SAN boot support

These sections provide information about SAN boot support.

Manual exclusion of devices from the SDD configuration

With certain technology levels of the AIX operating system, AIX supports fibre-channel boot capability for selected IBM System p and RS/6000 systems. This allows you to select fibre-channel devices as the boot device. However, a multipathing boot device is not supported. If you plan to select a device as a boot device, do not configure that device with multipath configuration. See the *Host System Attachment Guide* for the supported storage device for additional information.

The SDD configuration methods automatically excludes any devices from SDD configuration if these boot devices are the physical volumes of an active rootvg. If you require dual or multiple boot capabilities on a server and multiple operating systems are installed on multiple boot devices, use the **excludesddcfg** command to manually exclude all boot devices that belong to nonactive rootvg volume groups on the server or disk storage system devices that are going to be selected as boot devices.

SDD 1.3.3.9 (or later) allows you to manually exclude devices from the SDD configuration. See “*excludesddcfg*” on page 89 for the syntax of the **excludesddcfg** command. The maximum number of devices that can be excluded is 100.

You can also exclude any virtualization products devices from the SDD configuration with the `excludesddcfg` command.

Note:

1. Do not use the `excludesddcfg` command multiple times on the same logical device.
2. Fibre-channel boot capability is available for disk storage system only.

Replacing manually excluded devices in the SDD configuration

Use the following procedure to place manually excluded devices back in the SDD configuration.

Attention: Using this procedure will result in the loss of all data on these physical volumes. The data cannot be recovered.

1. If the excluded devices belong to an active volume group and file systems of that volume group are mounted, you need to perform one of the following actions:
 - a. Use the `umount` command to unmount all the file systems of the volume group and vary off the volume group.
 - b. Or, unmount all the file systems of the volume group and use the `reducevg` command to reduce that device from the volume group.
2. Issue `excludesddcfg -dl hdisk#` to remove this device from exclusion.
3. Issue `cfallvpath` configure methods to configure these new devices.
4. Issue `lsvpcfg` to verify that these devices are configured as SDD vpath devices.

SAN boot installation procedures

DS8000 supports single-path SAN boot on AIX 5.1 and later, while ESS and DS6000 support single-path SAN boot on AIX 5.1, AIX 5.2, and AIX 5.3. In order to successfully install the operating system on the single-path disk storage system devices in a multipath SDD environment, follow the provided procedures. The SAN boot procedures for AIX 5.1 are slightly different from the SAN boot procedures for AIX 5.2 and later.

Note: Because of current AIX boot device limitation, do not select more than 5 physical volumes (hdisks) for a boot device.

SAN boot installation procedure for AIX 5.1

Use this procedure for SAN boot installation for AIX 5.1:

1. Configure disk storage system devices to the AIX system; there should be only a single path per LUN. In other words, the AIX system should see only one hdisk configured per LUN.
2. Install the base OS on the selected disk storage system single-path devices.
3. Upgrade the base OS to latest technology level.
4. Install both SDD Host Attachment and SDD.
5. Reboot the AIX system.
6. Ensure that SDD vpath devices are configured correctly. Disk storage system devices (hdisks) should be configured as IBM 2105, IBM 2107, or IBM 1750 devices. Run the `datapath query device` command to verify that SDD vpath devices are configured with single-path and that the SDD vpath device policy is Single Path.

7. Issue **bosboot -aDd /dev/ipldevice**.
8. Connect additional host adapters and additional storage adapter to the fabric and configure multipath (multiple hdisks) per disk storage system LUN.
9. Reboot the AIX system.
10. Verify that SDD vpath devices are configured correctly with multipath per LUN.
11. Issue **bosboot -aDd /dev/ipldevice**.
12. Reboot the AIX system.

Note: The disk storage system's hdisk(s) name selected for the OS installation should not be changed throughout this procedure.

SAN boot installation procedure for AIX 5.2, AIX 5.3, and AIX 6.1

Use this procedure for SAN boot installation for AIX 5.2 and later releases:

1. Configure disk storage system devices to the AIX system; there should be only a single path per LUN. In other words, the AIX system should see only one hdisk configured per LUN.
2. Install the base operating system on the selected disk storage system single-path devices.
3. Upgrade the base operating system to the latest technology level.
4. Connect additional AIX host adapters and additional storage adapter to the fabric in order to configure multipath (multiple hdisks) per disk storage system LUN.
5. Install both SDD Host Attachment and SDD.
6. Reboot the AIX system.
7. Verify that SDD vpath devices are configured correctly with multipath per LUN. Disk Storage system devices (hdisks) should be configured as IBM 2105, IBM 2107, or IBM 1750 devices. Run the **datapath query device** command to verify that SDD vpath devices are configured with multiple paths and that the vpath device policy is Optimized.
8. The logical device names of the hdisks might not be configured in a continuous sequence because of the parallel configuration feature in AIX 5.2 and above. If that is the case, follow these additional steps to simplify future maintenance, before you create any SDD volume group and file systems.
 - a. Remove all hdisk logical devices names (rootvg will not be removed) and SDD vpath devices.
 - b. Reconfigure all the hdisk logical devices and SDD vpath devices with **cfgmgr** command, or reboot the AIX system.
 - c. Verify that all the logical device names of the hdisks (except rootvg) are configured in a continuous sequence.

Understanding SDD support for High Availability Cluster Multiprocessing

You can run SDD in concurrent and nonconcurrent multihost environments in which more than one host is attached to the same LUNs on a supported storage device. SDD supports High Availability Cluster Multiprocessing (HACMP) running on RS/6000 and IBM System p servers.

With SDD 1.4.0.0 (or later), there are no longer different SDD packages for HACMP running in concurrent and nonconcurrent modes. A single package (that corresponds to the AIX OS level) applies to HACMP running in different modes.

For SDD versions earlier than 1.4.0.0 but later than version 1.3.1.3, run the nonconcurrent version of SDD, if HACMP is running. For AIX 4.3, the nonconcurrent version of SDD would be `ibmSdd_433.rte`. For AIX 5.1, `ibmSdd_510nchacmp.rte` is the nonconcurrent version.

For SDD versions earlier than 1.3.1.3, see the corresponding User's Guide for HACMP support information. Table 9 provides a list of installation packages and the SDD versions (earlier than 1.4.0.0) and the HACMP modes that support them.

HACMP provides a reliable way for clustered RS/6000 and IBM System p servers that share disk resources to recover from server and disk failures. In an HACMP environment, each RS/6000 and IBM System p server in a cluster is a node. Each node has access to shared disk resources that other nodes access. When there is a failure, HACMP transfers ownership of shared disks and other resources based on how you define the resource takeover mode of the nodes in a cluster. This process is known as *node fallover* or *node fallback*. HACMP supports two modes of operation:

nonconcurrent

Only one node in a cluster is actively accessing shared disk resources while other nodes are standby.

concurrent

Multiple nodes in a cluster are actively accessing shared disk resources.

Table 9. Recommended SDD installation packages and supported HACMP modes for SDD versions earlier than SDD 1.4.0.0

Installation package	Version of SDD supported	HACMP mode supported
<code>ibmSdd_432.rte</code>	SDD 1.1.4 (SCSI only)	Concurrent
<code>ibmSdd_433.rte</code>	SDD 1.3.1.3 (or later) (SCSI and fibre channel)	Concurrent or nonconcurrent
<code>ibmSdd_510nchacmp.rte</code>	SDD 1.3.1.3 (or later) (SCSI and fibre channel)	Concurrent or nonconcurrent

Tip: If your SDD version is earlier than 1.4.0.0, and you use a *mix* of nonconcurrent and concurrent resource groups (such as *cascading and concurrent resource groups* or *rotating and concurrent resource groups*) with HACMP, you should use the nonconcurrent version of SDD.

Different storage systems or models might support different versions of HACMP. For information, see the interoperability matrix for your storage:

- www.ibm.com/systems/storage/disk/ess/
- www.ibm.com/systems/storage/disk/ds6000/
- www.ibm.com/systems/storage/disk/ds8000/
- www.ibm.com/systems/storage/software/virtualization/svc/

SDD supports RS/6000 and IBM System p servers connected to shared disks with SCSI adapters and drives as well as FCP adapters and drives. The kind of attachment support depends on the version of SDD that you have installed.

Table 10 on page 55 and Table 11 on page 55 summarize the software requirements to support HACMP v4.5.

You can use the command `instfix -ik IYxxxx`, where `xxxx` is the APAR number, to determine if APAR `xxxx` is installed on your system.

Table 10. Software support for HACMP 4.5 on AIX 4.3.3 (32-bit only), 5.1.0 (32-bit and 64-bit), 5.2.0 (32-bit and 64-bit)

SDD version and release level	HACMP 4.5 + APARs
devices.sdd.43.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)	Not applicable
devices.sdd.51.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)	<ul style="list-style-type: none"> • IY36938 • IY36933 • IY35735 • IY36951
devices.sdd.52.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)	<ul style="list-style-type: none"> • IY36938 • IY36933 • IY36782 • IY37744 • IY37746 • IY35810 • IY36951
<p>Note: For up-to-date APAR information for HACMP, go to the following website: www.ibm.com/support/us/en/</p>	

Table 11. Software support for HACMP 4.5 on AIX 5.1.0 (32-bit and 64-bit kernel)

SDD version and release level	HACMP 4.5 + APARs
ibmSdd_510nchacmp.rte installation package for SDD 1.3.1.3 (SCSI/FCP)	<ul style="list-style-type: none"> • IY36938 • IY36933 • IY35735 • IY36951
ibmSdd_510nchacmp.rte installation package for SDD 1.3.2.9 (to SDD 1.3.3.x) (SCSI/FCP)	<ul style="list-style-type: none"> • IY36938 • IY36933 • IY35735 • IY36951
<p>Note: For up-to-date APAR information for HACMP, go to the following website: www.ibm.com/support/us/en/</p>	

For HACMP v5.1, v5.2, v5.3, and v5.4 for AIX5L support information, go to the following website:

publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.hacmp.doc/hacmpbooks.html

For HACMP up-to-date APAR information, go to the following website:

www14.software.ibm.com/webapp/set2/sas/f/hacmp/download/aix53.html

SDD persistent reserve attributes

With SDD 1.4.0.0 or later, a single package (that corresponds to the AIX OS level) applies to HACMP running in both concurrent and nonconcurrent mode. In order to support HACMP in nonconcurrent mode with single-point-failure protection, the SDD installation packages implement the SCSI-3 Persistent Reserve command set.

The SDD installation packages have a new attribute under the pseudo-parent (dpo) that reflects whether the supported storage device supports the Persistent Reserve Command set. The attribute name is *persistent_resv*. If SDD detects that the storage device supports Persistent Reserve, the *persistent_resv* attribute is created in the CuAt ODM and the value is set to *yes*; otherwise this attribute exists only in the PdAt ODM and the value is set to *no* (default). You can use the following command to check the *persistent_resv* attribute after the SDD device configuration is complete:

```
lsattr -El dpo
```

If your host is attached to a supported storage device that supports Persistent Reserve, the output should look similar to the following output:

```
> lsattr -El dpo
  SDD_maxlun 1200 Maximum LUNS allowed True
  persistent_resv yes Subsystem Supports Persistent Reserve Command False
```

To check the persistent reserve key of a node that HACMP provides, enter the command:

```
odmget -q "name = ioaccess" CuAt
```

The output should look similar to this:

```
name = "ioaccess"
attribute = "preservekey"
value = "01043792"
type = "R"
generic = ""
rep = "s"
nls_index = 0
```

Preparation for importing volume groups under HACMP

Beginning with SDD 1.6.0.0, if the SDD vpath device is not reserved by another host and if there is pvid resided on the physical disk, the SDD configuration method will read the pvid from the physical disk and create the pvid attribute in the ODM database for all SDD vpath devices. Furthermore, the SDD configuration method will clean up supported storage devices' (hdisk) pvid from ODM database. With this feature, the host should have the pvid on the SDD vpath devices after an SDD vpath configuration, if a pvid exists on the physical disk (See Scenario 3 below). If no pvid exists on the physical disk, you will see the display as shown in Scenario 4 below. You should determine the scenario that matches your host and follow the procedures described for that scenario.

Before SDD 1.6.0.0, SDD does not automatically create the pvid attribute in the ODM database for each SDD vpath device. The AIX disk driver automatically creates the pvid attribute in the ODM database, if a pvid exists on the physical device. Because SDD versions earlier than 1.6.0.0 do not automatically create the pvid attribute in the ODM database for each SDD vpath device, the first time that you import a new SDD volume group to a new cluster node, you must import the volume group using hdisks as physical volumes. Next, run the hd2vp conversion script (see "SDD utility programs" on page 86) to convert the volume group's physical volumes from supported storage device hdisks to SDD vpath devices.

This conversion step not only creates pvid attributes for all SDD vpath devices that belong to that imported volume group, it also deletes the pvid attributes for the underlying hdisks for these SDD vpath devices. Later on, you can import and vary on the volume group directly from the SDD vpath devices. These special requirements apply to both concurrent and nonconcurrent volume groups.

Under certain conditions, the state of a physical device's pvid on a system is not always as expected. It is necessary to determine the state of a pvid as displayed by the `lspv` command, in order to select the appropriate import volume group action.

There are four scenarios:

Scenario 1. `lspv` displays pvids for both hdisks and vpath:

```
>lspv
hdisk1 003dfc10a11904fa None
hdisk2 003dfc10a11904fa None
vpath0 003dfc10a11904fa None
```

Scenario 2. `lspv` displays pvids for hdisks only:

```
>lspv
hdisk1 003dfc10a11904fa None
hdisk2 003dfc10a11904fa None
vpath0 none      None
```

For both Scenario 1 and Scenario 2, the volume group should be imported using the hdisk names and then converted using the `hd2vp` command:

```
>importvg -y vg_name -V major# hdisk1
>hd2vp vg_name
```

Scenario 3. `lspv` displays the pvid for vpath only:

```
>lspv
hdisk1 none      None
hdisk2 none      None
vpath0 003dfc10a11904fa None
```

For Scenario 3, the volume group should be imported using the vpath name:

```
>importvg -y vg_name -V major# vpath0
```

Scenario 4. `lspv` does not display the pvid on the hdisks or the vpath:

```
>lspv
hdisk1 none      None
hdisk2 none      None
vpath0 none      None
```

For Scenario 4, the pvid will need to be placed in the ODM for the SDD vpath devices and then the volume group can be imported using the vpath name:

```
>chdev -l vpath0 -a pv=yes
>importvg -y vg_name -V major# vpath0
```

Note: See "Importing volume groups with SDD" on page 78 for a detailed procedure for importing a volume group with the SDD devices.

HACMP RAID concurrent-mode volume groups and enhanced concurrent-capable volume groups

This section provides information about HACMP RAID concurrent-mode volume groups and enhanced concurrent-capable volume groups. This section also

provides instructions on the following procedures for both HACMP RAID concurrent-mode volume groups and enhanced concurrent-capable volume groups.

- Creating volume groups
- Importing volume groups
- Removing volume groups
- Extending volume groups
- Reducing volume groups
- Exporting volume groups

Starting with AIX51 TL02 and HACMP 4.4.1.4, you can create enhanced concurrent-capable volume groups with supported storage devices. HACMP supports both kinds of concurrent volume groups (HACMP RAID concurrent-mode volume groups and enhanced concurrent-capable volume groups).

This section describes the advantage of enhanced concurrent-capable volume groups in an HACMP environment. It also describes the different ways of creating two kinds of concurrent-capable volume groups. While there are different ways to create and vary on concurrent-capable volume groups, the instructions to export a volume group are always the same. See “Exporting HACMP RAID concurrent-mode volume groups” on page 63.

Note: For more information about HACMP RAID concurrent-mode volume groups, see the HACMP Administration Guide.

Creating HACMP RAID concurrent-mode volume groups

Complete the following steps to create an HACMP RAID concurrent-mode volume group:

Note: On each node in the cluster, issue the `lvlstmajor` command to determine the next common available major number (volume groups must be created with a major number that is available on all nodes).

1. Enter `smitty datapath_mkvg` at the command prompt.
2. A screen similar to the following is displayed. Enter the information appropriate for your environment. The following example shows how to create a concurrent access volume group `con_vg` on an SDD `vpath124`.

```

Add a Volume Group with Data Path Devices

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

          [Entry Fields]
VOLUME GROUP name                [con_vg]
Physical partition SIZE in megabytes
* PHYSICAL VOLUME names          [vpath124]
Force the creation of a volume group?      no
Activate volume group AUTOMATICALLY
  at system restart?              no
Volume group MAJOR NUMBER          [80]
Create VG Concurrent Capable?      no
Auto-varyon in Concurrent Mode?     no
Create a big VG format Volume Group?  no
LTG Size in kbytes                 128
  
```

Importing HACMP RAID concurrent-mode volume groups

When importing the volume group to other nodes in the cluster, you need to vary off the volume group on the node after it is created. You can import the volume group from either the SDD `vpath` device or the `hdisk` device, depending on the

pvid condition on the node to which the volume group is to be imported. Follow this procedure to import a volume group with SDD vpath device.

1. On the node where the volume group was originally created, you can get the pvid:

```

NODE VG ORIGINALLY CREATED ON
monkey> lspv | grep con_vg
vpath124 000900cf4939f79c con_vg
monkey>

```

2. Then **grep** the pvid on the other nodes using the **lspv | grep <pvid found in step 1>** and the **lsvpcfg** commands. There are three scenarios. Follow the procedure for the scenario that matches the pvid status of your host:

- a. If the pvid is on an SDD vpath device, the output of the **lspv | grep <pvid found in step 1>** and the **lsvpcfg** commands should look like the following example:

```

NODE VG BEING IMPORTED TO
zebra> lspv | grep 000900cf4939f79c
vpath124 000900cf4939f79c none
zebra>
zebra> lsvpcfg vpath124
vpath124 (Avail pv) 21B21411=hdisk126 (Avail) hdisk252 (Avail)

```

- 1) Enter **smitty importvg** at the command prompt.
- 2) A screen similar to the following example is displayed. Enter the information appropriate for your environment. The following example shows how to import an enhanced concurrent-capable volume group using the *con_vg* on an SDD vpath device *vpath3*:

```

*****
Import a Volume Group

Type or select values in the entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
VOLUME GROUP name          [con_vg]
PHYSICAL VOLUME names      [vpath124]
Volume Group MAJOR NUMBER  [80]
Make this VOLUME GROUP concurrent-capable?      no
Make default varyon of VOLUME GROUP concurrent? no
*****

```

- b. If the pvid is on hdisk devices, the output of the **lspv | grep <pvid found in step 1>** and the **lsvpcfg** commands should look like the following example:

```

NODE VG BEING IMPORTED TO
zebra> lspv | grep 000900cf4939f79c
hdisk126 000900cf4939f79c none
hdisk252 000900cf4939f79c none
zebra>
zebra> lsvpcfg | egrep -e 'hdisk126 ('
vpath124 (Avail) 21B21411=hdisk126 (Avail pv) hdisk252 (Avail pv)

```

- 1) Enter **smitty importvg** at the command prompt.
- 2) A screen similar to the following is displayed. Enter the information appropriate for your environment. The following example shows how to import an HACMP RAID concurrent-mode volume group using the *con_vg* on an SDD *hdisk126*:

```

*****
Import a Volume Group

Type or select values in the entry fields.
Press Enter AFTER making all desired changes.

```

```

[Entry Fields]
VOLUME GROUP name           [con_vg]
PHYSICAL VOLUME names       [hdisk126]
Volume Group MAJOR NUMBER   [80]
Make this VOLUME GROUP concurrent-capable?  no
Make default varyon of VOLUME GROUP concurrent?  no
*****

```

- 3) After importing volume groups have been completed, issue the **lsvpcfg** command again to verify the state of the vpath.

```

zebra> lsvpcfg | egrep -e 'hdisk126 ('
vpath124 (Avail) 21B21411=hdisk126 (Avail pv con_vg) hdisk252 (Avail pv con_vg)

```

- 4) Enter the **hd2vp** command against the volume group to convert the pvid from hdisk devices to SDD vpath devices:

```

zebra> hd2vp con_vg
zebra> lsvpcfg | egrep -e 'hdisk126 ('
vpath124 (Avail pv con_vg) 21B21411=hdisk126 (Avail) hdisk252 (Avail)

```

- c. If there is no pvid on either hdisk or SDD vpath device, the output of the **lspv | grep <pvid found in step 1>** and the **lsvpcfg** commands should look like the following example:

```

NODE VG BEING IMPORTED TO
zebra> lspv | grep 000900cf4939f79c
zebra>

```

- 1) Issue the **chdev -l vpathX -a pv=yes** command to retrieve the pvid value.
- 2) There is a possibility that the SDD vpath device might be different for each host. Verify that the serial numbers (in this example, it is 21B21411) following the SDD vpath device names on each node are identical. To determine a matching serial number on both nodes, run the **lsvpcfg** command on both nodes.

```

monkey> lsvpcfg
vpath122 (Avail) 21921411=hdisk255 (Avail) hdisk259 (Avail)
vpath123 (Avail) 21A21411=hdisk256 (Avail) hdisk260 (Avail)
vpath124 (Avail pv con_vg) 21B21411=hdisk127 (Avail) hdisk253 (Avail)
monkey>

zebra> lsvpcfg | egrep -e '21B21411
vpath124 (Avail) 21B21411=hdisk126 (Avail) hdisk252 (Avail)
zebra>

```

Note: You should also verify that the volume group is not varied on for any of the nodes in the cluster prior to attempting retrieval of the pvid.

- 3) Enter **smitty importvg** at the command prompt.
- 4) A screen similar to the following is displayed. Enter the information appropriate for your environment. The following example shows how to import an HACMP RAID concurrent-mode volume group using the *con_vg* on an SDD vpath device *vpath124*.

```

*****
Import a Volume Group

Type or select values in the entry fields.
Press Enter AFTER making all desired changes.

```

```

[Entry Fields]
VOLUME GROUP name           [con_vg]
PHYSICAL VOLUME names       [vpath124]
Volume Group MAJOR NUMBER   [80]

```

```

Make this VOLUME GROUP concurrent-capable? no
Make default varyon of VOLUME GROUP concurrent? no
*****

```

3. After importing volume groups has been completed, issue the `lsvpcfg` command again to verify the state of the SDD vpath device.

```

zebra> lsvpcfg vpath124
vpath124 (Avail pv con_vg) 21B21411=hdisk126 (Avail) hdisk252 (Avail)

```

Attention: When any of these HACMP RAID concurrent-mode volume groups are imported to the other nodes, it is important that they are *not* set for `autovaryon`. This will cause errors when attempting to synchronize the HACMP cluster. When the concurrent access volume groups are not set to `autovaryon`, a special option flag `-u` is required when issuing the `varyonvg` command to make them concurrent-accessible across all the cluster nodes.

Use the `lsvg vname` command to check the value of `autovaryon`.

Use the `chvg -an vname` command to set `autovaryon` to `FALSE`.

Removing HACMP RAID concurrent-mode volume groups

Note:

1. Removing an HACMP RAID concurrent-mode volume group can be accomplished by exporting volume groups, or by following the procedure below.
2. These steps need to be run on all nodes.

Complete the following steps to remove an HACMP RAID concurrent-mode volume group:

1. Ensure that the volume group is varied on.
2. Enter `smitty vg` at the command prompt.
3. Select **Remove a Volume Group** from the displayed menu.

Note: A screen similar to the following example is displayed. Enter the information appropriate for your environment. The following example shows how to remove an HACMP RAID concurrent-mode volume group using the `con_vg` volume group.

```

*****
Remove a Volume Group

Type or select values in the entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
VOLUME GROUP name          [con_vg]
*****

```

Extending HACMP RAID concurrent-mode volume groups

Complete the following steps to extend an HACMP RAID concurrent-mode volume group:

1. Vary off the HACMP RAID concurrent-mode volume group to be extended on all nodes.
2. Enter `smitty datapath_extendvg` at the command prompt of one of the nodes.

3. A screen similar to the following example is displayed. Enter the information appropriate for your environment. The following example shows how to extend an HACMP RAID concurrent-mode volume group using the *con_vg* on an SDD *vpath2*.

```
*****
Add a Datapath Physical Volume to a Volume Group

Type or select values in the entry fields.
Press Enter AFTER making all desired changes.
```

```
[Entry Fields]
VOLUME GROUP name          [con_vg]
PHYSICAL VOLUME names     [vpath2]
*****
```

4. Vary off the volume group after extending it on the current node.
5. For all the nodes sharing *con_vg*, do the following:
 - a. Enter the **chdev -l vpath2 -a pv=yes** command to obtain the pvid for this vpath on the other host.
 - b. Verify that the pvid exists by issuing the **lspv** command.
 - c. Enter **importvg -L con_vg vpath2** to import the volume group again.
 - d. Verify that *con_vg* has the extended vpath included by using the **lspv** command.

Reducing HACMP RAID concurrent-mode volume groups

Complete the following steps to reduce an HACMP RAID concurrent-mode volume group:

1. Vary off the HACMP RAID concurrent-mode volume group to be reduced on all nodes.
2. Enter **smitty vg** at the command prompt.
3. Select **Set Characteristics of a Volume Group** from the displayed menu.
4. Select **Remove a Physical Volume from a Volume Group** from the displayed menu.
5. A screen similar to the following example is displayed. Enter the information appropriate for your environment. The following example shows how to reduce an HACMP RAID concurrent-mode volume group using the *con_vg* on an SDD *vpath1*. Assume that *con_vg* originally has *vpath0* and *vpath1* as its physical volumes.

```
*****
Remove a Physical Volume from a Volume Group

Type or select values in the entry fields.
Press Enter AFTER making all desired changes.
```

```
[Entry Fields]
VOLUME GROUP name          [con_vg]
PHYSICAL VOLUME names     [vpath1]
FORCE deallocation of all partitions  yes
*****
```

6. Vary off the volume group after reducing it on the current node.
7. For all the nodes sharing *con_vg*, do the following:
 - a. Enter **exportvg con_vg** at the command prompt.
 - b. Enter **smitty importvg** at the command prompt.

- c. A screen similar to the following is displayed. Enter the information appropriate for your environment.

```

*****
                          Import a Volume Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
VOLUME GROUP name                [con_vg]
PHYSICAL VOLUME name             [vpath0]
+
Volume Group MAJOR NUMBER        [45]
+#
Make this VG Concurrent Capable?  No      +
Make default varyon of VG Concurrent? no     +
*****

```

- d. Verify that `con_vg` has the `vpath` reduced by using the `lspv` command.

Exporting HACMP RAID concurrent-mode volume groups

To export an HACMP RAID concurrent-mode volume group, enter `exportvg` at the command prompt.

Note:

1. To export `con_vg`, use the `exportvg con_vg` command.
2. Before exporting an HACMP RAID concurrent-mode volume group, make sure the volume group is varied off.

Enhanced concurrent-capable volume groups

With the AIX51 TL02 and HACMP 4.4.1.4 environments, enhanced concurrent mode is supported with both 32-bit and 64-bit kernels. The advantage of this mode is that after you create an enhanced concurrent-capable volume group on multiple nodes, the changes made to the logical volume or volume group structures on one node (for example, extending or reducing a volume group), are propagated to all other nodes. Also, the Logical Volume Manager (LVM) configuration files are updated on all nodes.

The following sections provide information and instructions on the operating actions that you can perform.

For more detailed information on enhanced concurrent-capable volume groups, see “Supporting enhanced concurrent mode in an HACMP environment” on page 66.

Creating enhanced concurrent-capable volume groups:

Complete the following steps to create enhanced concurrent-capable volume groups:

Note: On each node in the cluster, issue the `lvlstmajor` command to determine the next available major number. The volume groups must be created with a major number that is available on all nodes. The following listing is an example:

```

dollar>lvfstmajor
41,54..58,60..62,67,78...
monkey>lvfstmajor
39..49,55,58,67,80...
zebra>lvfstmajor
40..49,55,58..61,67,78...

```

From this listing, the next common available major number can be selected (41, 55, 58, 61, 67, 68, 80, ...). However, if multiple volume groups are going to be created, the user might begin with the highest available (80) and increase by increments from there.

1. Enter **smitty datapath_mkvg** at the command prompt.
2. A screen similar to the following example is displayed. Enter the information appropriate for your environment. The following example shows how to create an enhanced concurrent-capable volume group using the *con_vg* on an SDD *vpath0*.

```

Add a Volume Group with Data Path Devices

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

VOLUME GROUP name                [Entry Fields]
Physical partition SIZE in megabytes [con_vg]
* PHYSICAL VOLUME names            [vpath0]
Force the creation of a volume group? no
Activate volume group AUTOMATICALLY at system restart? no
Volume group MAJOR NUMBER          [80]
Create VG Concurrent Capable?      yes
Auto-varyon in Concurrent Mode?    no
Create a big VG format Volume Group? no
LTG Size in kbytes                  128

```

Importing enhanced concurrent-capable volume groups:

Complete the following steps to import enhanced concurrent-capable volume groups.

Before importing enhanced concurrent-capable volume groups on SDD *vpath* devices, issue the **lspv** command to make sure there is *pvid* on the SDD *vpath* device. If *pvid* is not displayed, import the volume group on one of the *hdisks* that belongs to the SDD *vpath* device. Enter **hd2vp** to convert the volume group to SDD *vpath* devices.

If the *hdisks* do not have a *pvid*, run **chdev -l hdiskX -a pv=yes** to recover it. To verify that *pvid* now exists, run the **lspv** command against the *hdisk*. This method can also be used when attempting to obtain a *pvid* on an SDD *vpath* device.

Verify that the volume group is not varied on for any of the nodes in the cluster prior to attempting to retrieve the *pvid*.

Enter **smitty importvg** at the command prompt. A screen similar to the following example is displayed. Enter the information appropriate to your environment. The following example shows how to import an enhanced concurrent-capable volume group using the *con_vg* on SDD *vpath* device *vpath3*.

```

*****
Import a Volume Group

Type or select values in the entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
VOLUME GROUP name                [con_vg]
PHYSICAL VOLUME names            [vpath3]
Volume Group MAJOR NUMBER        [45]

```

```

Make this VOLUME GROUP concurrent-capable?          yes
Make default varyon of VOLUME GROUP concurrent?    no
*****

```

Note: The major number identified must be the same one used when the volume group was first created.

Extending enhanced concurrent-capable volume groups:

Note: Before attempting the extend of the concurrent volume group, ensure that pvids exist on the SDD vpath device/hdisks on all nodes in the cluster.

Complete the following steps to extend an enhanced concurrent-capable volume group:

1. Enter **smitty datapath_extendvg** at the command prompt.
2. A screen similar to the following is displayed. Enter the information appropriate for your environment. The following example shows how to extend an enhanced concurrent-capable volume group using the *con_vg* on SDD vpath device *vpath2*.

```

*****
Add a Datapath Physical Volume to a Volume Group

Type or select values in the entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
VOLUME GROUP name          [con_vg]
PHYSICAL VOLUME names      [vpath2]
*****

```

Note: Verify that extending of enhanced concurrent-capable volume groups worked on the particular node and that all changes were propagated to all other nodes in the cluster using the **lsvpcfg** command.

Reducing enhanced concurrent-capable volume groups:

Complete the following steps to reduce an enhanced concurrent-capable volume group:

1. Enter **smitty vg** at the command prompt.
2. Select **Set Characteristics of a Volume Group** from the displayed menu.
3. Select **Remove a Physical Volume from a Volume Group** from the displayed menu.
4. A screen similar to the following is displayed. Enter the information appropriate for your environment. The following example shows how to reduce an enhanced concurrent-capable volume group using the *con_vg* on SDD vpath device *vpath2*.

```

*****
Remove a Physical Volume from a Volume Group

Type or select values in the entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
VOLUME GROUP name          [con_vg]

```

```

PHYSICAL VOLUME names          [vpath2]
FORCE deallocation of all partitions      yes
*****

```

Note: Verify that reducing of volume groups worked on the particular node and that all changes were propagated to all other nodes in the cluster using the `lsvpcfg` command.

Recovering paths that are lost during HACMP node failover that is caused when a system locks up

Typically, if an active node locks up, HACMP transfers ownership of shared disks and other resources through a process known as *node failover*. Certain situations, such as a loose or disconnected SCSI or fibre-channel-adapter card, can cause your SDD vpath devices to lose one or more underlying paths after the failed node is restarted. Complete the following steps to recover these paths:

- Make sure the issue that is causing lost paths is fixed. Then run the `cfgmgr` command to configure all the underlying paths (hdisks) to *Available* state.
- Enter the `addpaths` command to add the lost paths back to the SDD devices.

If your SDD vpath devices have lost one or more underlying paths that belong to an active volume group, you can use either the Add Paths to Available Data Path Devices SMIT panel or run the `addpaths` command from the AIX command line to recover the lost paths. Go to “Dynamically adding paths to SDD vpath devices” on page 46 for more information about the `addpaths` command.

Note: Running the `cfgmgr` command while the SDD vpath devices are in the *Available* state will not recover the lost paths; you must run the `addpaths` command to recover the lost paths.

Supporting enhanced concurrent mode in an HACMP environment

To run HACMP in this enhanced concurrent mode, you need:

- ESCRM feature of HACMP
- `bos.clvm.enh` and `bos.rte.lvm` filesets installed at level 5.1.0.10 (or later) on all the nodes

SDD 1.3.2.9 (or later) provides the updated version of `mkvg4vp` and smit panel for the user to create enhanced concurrent-capable volume groups. To create enhanced concurrent-capable volume groups from the command line, the user needs to turn on the `-c` (in 32-bit kernel) or the `-C` (in 64-bit kernel) option for the `mkvg4vp` command. To create enhanced concurrent-capable volume groups from the smit panel, set **Create Volume Group concurrent-capable?** to *yes*. Both ways will leave the enhanced concurrent-capable volume group in varied-off mode. Import this concurrent volume group to all other nodes and add the volume group into the HACMP concurrent resource group, and then start the HACMP cluster. The volume group will be varied-on by HACMP. After the changes are made to one volume group on one node, all changes are automatically propagated to the other nodes.

For more detailed information and instructions on creating, removing, reducing, importing, and exporting enhanced concurrent-capable volume groups, see “Enhanced concurrent-capable volume groups” on page 63.

SDD server daemon

The SDD server (sddsrv) is an integrated component of SDD 1.3.2.9 (or later). This component consists of a UNIX application daemon that is installed in addition to the SDD device driver. See Chapter 9, “Using the SDDDSM or the SDDPCM daemon,” on page 365 for more information about sddsrv.

Attention: Running sddsrv will activate several AIX Fibre Channel Protocol or adapter driver problems on AIX 4.3.3 and 5.1.0. One of the problems in the AIX Fibre Channel Protocol driver is that internal resources can be leaked. You will experience this as a performance degradation that grows worse over time. Performance can be restored by unconfiguring and reconfiguring the fibre-channel adapter or by rebooting the system. AIX users with Fibre Channel Protocol support and the SDD server daemon installed should apply the PTFs listed in “PTFs for APARs on AIX with Fibre Channel and the SDD server” on page 68.

Verifying if the SDD server has started

After you have installed SDD, verify if the SDD server (sddsrv) has automatically started by entering `lssrc -s sddsrv`.

If the SDD server (sddsrv) has automatically started, the output from `lssrc -s sddsrv` command looks like this:

```
Subsystem  GROUP  PID Status
sddsrv      NNN  Active
```

where *NNN* is the process ID number.

The status of sddsrv should be *Active* if the SDD server has automatically started.

If the SDD server has *not* started, the status will be *Inoperative*. Go to “Starting the SDD server manually” to proceed.

Note: During OS installations and migrations, the following command could be added to `/etc/inittab`:

```
install_assist:2:wait:/usr/sbin/install_assist </dev/console>/dev/console 2>&1
```

Because this command runs in the foreground, it will prevent all the subsequent commands in the script from starting. If sddsrv happens to be behind this line, sddsrv will not run after system reboot. You should check `/etc/inittab` during OS installations or migrations and comment out this line.

Starting the SDD server manually

If the SDD server did not start automatically after you performed the SDD installation, you can start sddsrv by entering `startsrc -s sddsrv`.

Go to “Verifying if the SDD server has started” to verify that the SDD server started successfully.

Changing to a different port number for the SDD server

To change to a different port number for the SDD server, see the instructions in “Changing the sddsrv or pcmsrv TCP/IP port number” on page 369.

Stopping the SDD server

You can temporarily disable sddsrv by entering the command **stopsrc -s sddsrv**. This will stop the current version of sddsrv, but sddsrv will start again if the system is restarted.

If sddsrv is stopped, the feature that automatically recovers failed paths (DEAD or CLOSE_DEAD path) is disabled. During the concurrent storage bay quiesce/resume process, you must manually recover the adapter or paths after the quiesce/resume is completed on one bay, and before the quiesce/resume starts on the next bay. Without doing so, the application might fail. See “datapath set device path” on page 390 for more information.

Replacing the SDD server with a stand-alone version

You can also choose to replace the current version of sddsrv with a stand-alone version by performing the following steps:

1. Run **stopsrc -s sddsrv** to stop the current version of sddsrv.
2. Verify that the SDD server has stopped. See “Verifying if the SDD server has started” on page 67 and the status should be *inoperative*.
3. Comment out following line in the system /etc/inittab table:

```
srv:2:wait:/usr/bin/startsrc -s sddsrv > /dev/null 2>&1
```

4. Add following line to the system /etc/inittab table:

```
srv:2:wait:/usr/bin/startsrc -a s0 -s sddsrv > /dev/null 2>&1
```

5. Run **startsrc -a s0 -s sddsrv** to start a stand-alone version of sddsrv.

Starting sddsrv with the s0 flag does not provide path health check or path reclamation functions. You should manually recover paths by using the **datapath** command. See “datapath set device path” on page 390 for more information.

PTFs for APARs on AIX with Fibre Channel and the SDD server

If you have fibre-channel support and the SDD server daemon running, PTFs for the APARs shown in Table 12 must be applied to your AIX servers in order to avoid a performance degradation.

Table 12. PTFs for APARs on AIX with fibre-channel support and the SDD server daemon running

AIX version	APAR	PTF
AIX 5.1	IY32325 (available in either of devices.pci.df1000f7.com 5.1.0.28 or 5.1.0.35)	U476971 U482718
AIX 5.1	IY37437 (available in devices.pci.df1000f7.com 5.1.0.36)	U483680
AIX 4.3.3	IY35177 (available in devices.pci.df1000f7.com 4.3.3.84)	U483803
AIX 4.3.3	IY37841 (available in devices.pci.df1000f7.com 4.3.3.86)	U484723

If you experience a degradation in performance, you should disable sddsrvc until the PTFs for these APARs can be installed. After the PTFs for these APARs are installed, you should re-enable sddsrvc. If you are running IBM TotalStorage Expert, see the “Replacing the SDD server with a stand-alone version” on page 68. Otherwise, see “Stopping the SDD server” on page 68.

Dynamically changing the SDD path-selection policy algorithm

SDD 1.3.3.9 (or later) supports dynamically changing the SDD devices path-selection policy. The following path-selection policies are supported:

failover only (fo)

All I/O operations for the device are sent to the same (preferred) path until the path fails because of I/O errors. Then an alternate path is chosen for subsequent I/O operations. If there are multiple (preferred) paths on multiple adapters, I/O operation on each adapter will *not* be balanced among the adapters based on the load of each adapter.

load balancing (lb)

The path to use for an I/O operation is chosen by estimating the load on the adapter to which each path is attached. The load is a function of the number of I/O operations currently in process. If multiple paths have the same load, a path is chosen at random from those paths. Load-balancing mode also incorporates failover protection.

Note: The load-balancing policy is also known as the *optimized policy*.

load balancing sequential (lbs)

This policy is the same as the load-balancing policy with optimization for sequential I/O.

round robin (rr)

The path to use for each I/O operation is chosen at random from paths that were not used for the last I/O operation. If a device has only two paths, SDD alternates between the two.

round robin sequential (rrs)

This policy is the same as the round-robin policy with optimization for sequential I/O.

default (df)

The policy is set to the default policy, which is load balancing.

The path-selection policy is set at the SDD device level. The default path-selection policy for an SDD device is load balancing.

Before changing the path-selection policy, determine the active attributes for the SDD device. Enter the `lsattr -El vpathN` command, where *N* represents the vpath number. Press **Enter**. The output should look similar to this:

```
# lsattr -El vpath0
active_hdisk hdisk66/13AB2ZA1020/fscsi3      Active hdisk      False
active_hdisk hdisk2/13AB2ZA1020/fscsi2      Active hdisk      False
active_hdisk hdisk34/13AB2ZA1020/fscsi2     Active hdisk      False
active_hdisk hdisk98/13AB2ZA1020/fscsi3     Active hdisk      False
policy       df                                           Scheduling Policy True
pvid        0005f9fdcda4417d00000000000000000 Physical volume identifier False
qdepth_enable no                               Queue Depth Control True
reserve_policy PR_exclusive                    Reserve Policy     True
serial_number 13AB2ZA1020                      LUN serial number False
unique_id    yes                               Device Unique Identification False
```

datapath set device policy command

Use the **datapath set device policy** command to change the SDD path-selection policy dynamically:

Note: You can enter the **datapath set device *N* policy** command to change the policy dynamically associated with vpaths in either *Close* or *Open* state.

See “datapath set device policy” on page 389 for more information about the **datapath set device policy** command.

Dynamically opening an invalid or close_dead path

With SDD 1.3.2.9 (or later), you can issue the **datapath open path** command to dynamically open a path that is in an INVALID or CLOSE_DEAD state if the SDD vpath device it belongs to is in the OPEN state. You can use this command even when the I/O is actively running. See “datapath open device path” on page 374 in Chapter 10, “Using the datapath commands,” on page 371 for more information.

Fibre-channel dynamic device tracking for AIX 5.20 TL1 (and later)

This section applies only to AIX 5.20 TL1 and later releases.

Beginning with AIX 5.20 TL1, the AIX fibre-channel driver supports fibre-channel dynamic device tracking. This enables the dynamic changing of fibre-channel cable connections on switch ports or on supported storage ports without unconfiguring and reconfiguring hdisk and SDD vpath devices.

With dynamic tracking enabled, the fibre-channel adapter detects the change of the device's fibre-channel node port ID. It reroutes the traffic that is destined for that device to the new worldwide port name (WWPN) while the device is still online.

SDD 1.5.0.0 and later support this feature. SDD 1.6.0.0 and later support all disk storage system devices. This feature allows for the following scenarios to occur without I/O failure:

1. Combine two switches in two SANs into one SAN by connecting switches with cable and cascading switches within 15 seconds.
2. Change connection to another switch port; the disconnected cable must be reconnected within 15 seconds.
3. Swap switch ports of two cables on the SAN; the disconnected cable must be reconnected within 15 seconds. The switch ports must be in the same zone on the same switch.
4. Swap ports of two cables on disk storage system; the disconnected cable must be reconnected within 15 seconds.

Note: This 15 seconds includes the time to bring up the fibre channel link after you reconnect the cables. Thus the actual time that you can leave the cable disconnected is less than 15 seconds. For disk storage systems, it takes approximately 5 seconds to bring up the fibre channel link after the fibre channel cables are reconnected.

By default, dynamic tracking is disabled. Use the following procedure to enable dynamic tracking:

1. Issue the **rmdev -l fscsiX -R** for all adapters on your system to change all the child devices of fscsiX on your system to the defined state.

2. Issue the `chdev -l fscsiX -a dyntrk=yes` command for all adapters on your system.
3. Run `cfgmgr` to reconfigure all devices back to the available state.

To use Fibre-channel Dynamic Device Tracking, you need the following fibre-channel device driver PTFs applied to your system:

- U486457.bff (This is a prerequisite PTF.)
- U486473.bff (This is a prerequisite PTF.)
- U488821.bff
- U488808.bff

After applying the PTFs listed above, use the `lspp` command to ensure that the files `devices.fcp.disk.rte` and `devices.pci.df1000f7.com` are at level 5.2.0.14 or later.

Note: Fibre-channel device dynamic tracking does not support the following case:

The port change on the supported storage devices where a cable is moved from one adapter to another free, previously unseen adapter on the disk storage system. The World Wide Port Name will be different for that previously unseen adapter, and tracking will not be possible. The World Wide Port Name is a static identifier of a remote port.

Understanding SDD 1.3.2.9 (or later) support for single-path configuration for supported storage devices

SDD 1.3.2.9 (or later) does not support concurrent download of licensed machine code in single-path mode.

SDD does support single-path SCSI or fibre-channel connection from your AIX host system to supported storage devices. It is possible to create a volume group or an SDD vpath device with only a single path. However, because SDD cannot provide single-point-failure protection and load balancing with a single-path configuration, you should not use a single-path configuration.

Tip: It is also possible to change from single-path to multipath configuration by using the `addpaths` command. For more information about the `addpaths` command, go to “Dynamically adding paths to SDD vpath devices” on page 46.

Understanding the persistent reserve issue when migrating from SDD to non-SDD volume groups after a system reboot

There is an issue with migrating from SDD to non-SDD volume groups after a system reboot. This issue only occurs if the SDD volume group was varied on prior to the system reboot and `auto varyon` was not set when the volume group was created. After the system reboot, the volume group will not be varied on.

The command to migrate from SDD to non-SDD volume group (`vp2hd`) will succeed, but a subsequent command to vary on the volume group will fail. This is because during the reboot, the persistent reserve on the physical volume of the volume group was not released, so when you vary on the volume group, the command will do a SCSI-2 reserve and fail with a reservation conflict.

There are two ways to avoid this issue.

1. Unmount the filesystems and vary off the volume groups before rebooting the system.
2. Issue `lquerypr -Vh /dev/vpathX` on the physical LUN before varying on volume groups after the system reboot. If the LUN is reserved by the current host, release the reserve by issuing `lquerypr -Vrh /dev/vpathX` command. After successful processing, you can vary on the volume group successfully.

Managing secondary-system paging space

SDD 1.3.2.6 (or later) supports secondary-system paging on *multipath* fibre-channel SDD vpath device from an AIX 4.3.3, AIX 5.1.0, AIX 5.2, AIX 5.3, or AIX 6.1 host system to a supported storage device.

SDD supports secondary-system paging on supported storage devices. The benefits are multipathing to your paging spaces. All the same commands for hdisk-based volume groups apply to using vpath-based volume groups for paging spaces.

The following sections provide information about managing secondary-system paging space.

Note: AIX does not recommend moving the primary paging space out of rootvg. Doing so might mean that *no* paging space is available during the system startup, which can result in poor startup performance. *Do not* redefine your primary paging space using SDD vpath devices.

Listing paging spaces

You can list paging spaces by entering:

```
lsps -a
```

Adding a paging space

You can add a paging space by entering:

```
mkps -a -n -sNN vg
```

The `mkps` command recognizes the following options and arguments:

- `-a` Makes the new paging space available at all system restarts.
- `-n` Activates the new paging space immediately.
- `-sNN` Accepts the number of logical partitions (*NN*) to allocate to the new paging space.
- vg* The volume group name in which a paging logical volume is to be created.

Removing a paging space

You can remove a specified secondary paging space that is not active.

For example, to remove paging space PS01, enter:

```
rmpps PS01
```

Verifying load-balancing and failover protection

SDD provides load-balancing and failover protection for AIX applications and for the LVM when SDD vpath devices are used. These devices must have a minimum of two paths to a physical LUN for failover protection to exist.

Displaying the supported storage device SDD vpath device configuration

To provide failover protection, an SDD vpath device must have a minimum of two paths. Both the SDD vpath device and the hdisk devices must be in the *Available* state. In the following example, vpath0, vpath1, and vpath2 all have a single path and, therefore, will not provide failover protection because there is no alternate path to the LUN. The other SDD vpath devices have two paths and, therefore, can provide failover protection.

To display which supported storage device SDD vpath devices are available to provide failover protection, use either the Display Data Path Device Configuration SMIT panel, or run the `lsvpcfg` command. Complete the following steps to use SMIT:

1. Enter **smitty device** from your desktop window. The Devices panel is displayed.
2. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
3. Select **Display Data Path Device Configuration** and press **Enter**.
4. To display the state (either Defined or Available) of all SDD vpath devices and the paths to each device, select all devices for **Select Query Option**, leave Device Name/ Device Model blank and press **Enter**.

You will see an output similar to the following example:

```
vpath0 (Avail pv vpathvg) 018FA067 = hdisk1 (Avail )
vpath1 (Avail ) 019FA067= hdisk2 (Avail )
vpath2 (Avail ) 01AFA067 = hdisk3 (Avail )
vpath3 (Avail ) 01BFA067 = hdisk4 (Avail ) hdisk27 (Avail )
vpath4 (Avail ) 01CFA067 = hdisk5 (Avail ) hdisk28 (Avail )
vpath5 (Avail ) 01DFA067 = hdisk6 (Avail ) hdisk29 (Avail )
vpath6 (Avail ) 01EFA067 = hdisk7 (Avail ) hdisk30 (Avail )
vpath7 (Avail ) 01FFA067 = hdisk8 (Avail ) hdisk31 (Avail )
vpath8 (Avail ) 020FA067 = hdisk9 (Avail ) hdisk32 (Avail )
vpath9 (Avail pv vpathvg) 02BFA067 = hdisk20 (Avail ) hdisk44 (Avail )
vpath10 (Avail pv vpathvg) 02CFA067 = hdisk21 (Avail ) hdisk45 (Avail )
vpath11 (Avail pv vpathvg) 02DFA067 = hdisk22 (Avail ) hdisk46 (Avail )
vpath12 (Avail pv vpathvg) 02EFA067 = hdisk23 (Avail ) hdisk47 (Avail )
vpath13 (Avail pv vpathvg) 02FFA067 = hdisk24 (Avail ) hdisk48 (Avail )
```

The following information is displayed:

- The name of each SDD vpath device, such as vpath1.
- The configuration state of the SDD vpath device. It is either *Defined* or *Available*. There is no failover protection if only one path is in the *Available* state. At least two paths to each SDD vpath device must be in the *Available* state to have failover protection.

Attention: The configuration state also indicates whether the SDD vpath device is defined to AIX as a physical volume (pv flag). If pv is displayed for *both* SDD vpath devices and the hdisk devices that it is comprised of, you might not have failover protection. Enter the `dpovgfix` command to fix this problem.

- The name of the volume group to which the device belongs, such as vpathvg.

- The unit serial number of the supported storage device LUN, such as 019FA067.
- The names of the AIX disk devices that comprise the SDD vpath devices, their configuration states, and the physical volume states. See “lsvpcfg” on page 87 for information about the **lsvpcfg** command.

You can also use the **datapath** command to display information about an SDD vpath device. This command displays the number of paths to the device. For example, the **datapath query device 10** command might produce this output:

```

DEV#: 10 DEVICE NAME: vpath10 TYPE: 2105B09 POLICY: Optimized
SERIAL: 02CFA067
-----
Path#      Adapter/Hard Disk  State   Mode    Select  Errors
   0         scsi6/hdisk21     OPEN   NORMAL    44       0
   1         scsi5/hdisk45     OPEN   NORMAL    43       0

```

The sample output shows that device vpath10 has two paths and both are operational. See “datapath query device” on page 378 for more information about the **datapath query device** command.

Configuring volume groups for failover protection

You can create a volume group with SDD vpath devices using the Volume Groups SMIT panel. Choose the SDD vpath devices that have failover protection for the volume group.

It is possible to create a volume group that has only a single path (see 73) and then add paths later by reconfiguring the supported storage device. (See “Dynamically adding paths to SDD vpath devices” on page 46 for information about adding paths to an SDD device.) However, an SDD volume group does not have failover protection if any of its physical volumes has only a single path.

Complete the following steps to create a new volume group with SDD vpath devices:

1. Enter **smitty** at the AIX command prompt. The System Management Interface Tool (SMIT) is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Group** and press **Enter**. The **Volume Groups** panel is displayed.
5. Select **Add a Volume Group with Data Path Devices** and press **Enter**.

Note: Press **F4** while highlighting the **PHYSICAL VOLUME names** field to list all the available SDD vpath devices.

If you use a script file to create a volume group with SDD vpath devices, you must modify your script file and replace the **mkvg** command with the **mkvg4vp** command.

All the functions that apply to a regular volume group also apply to an SDD volume group. Use SMIT to create a logical volume (mirrored, striped, or compressed) or a file system (mirrored, striped, or compressed) on an SDD volume group.

After you create the volume group, AIX creates the SDD vpath device as a physical volume (pv). In the output shown on page 73, vpath9 through vpath13 are included in a volume group and they become physical volumes. To list all the physical volumes known to AIX, use the **lspv** command. Any SDD vpath devices that were created into physical volumes are included in the output similar to the following output:

```

hdisk0      0001926922c706b2  rootvg
hdisk1      none                None
...
hdisk10     none                None
hdisk11     00000000e7f5c88a  None
...
hdisk48     none                None
hdisk49     00000000e7f5c88a  None
vpath0      00019269aa5bc858  None
vpath1      none                None
vpath2      none                None
vpath3      none                None
vpath4      none                None
vpath5      none                None
vpath6      none                None
vpath7      none                None
vpath8      none                None
vpath9      00019269aa5bbadd  vpathvg
vpath10     00019269aa5bc4dc  vpathvg
vpath11     00019269aa5bc670  vpathvg
vpath12     000192697f9fd2d3  vpathvg
vpath13     000192697f9fde04  vpathvg

```

To display the devices that comprise a volume group, enter the **lsvg -p vg-name** command. For example, the **lsvg -p vpathvg** command might produce the following output:

PV_NAME	PV STATE	TOTAL PPs	FREE PPs	FREE DISTRIBUTION
vpath9	active	29	4	00..00..00..00..04
vpath10	active	29	4	00..00..00..00..04
vpath11	active	29	4	00..00..00..00..04
vpath12	active	29	4	00..00..00..00..04
vpath13	active	29	28	06..05..05..06..06

The example output indicates that the **vpathvg** volume group uses physical volumes vpath9 through vpath13.

Losing failover protection

AIX can create volume groups only from or SDD vpath devices that are physical volumes. If a volume group is created using a device that is not a physical volume, AIX makes it a physical volume as part of the procedure of creating the volume group. A physical volume has a physical volume identifier (pvid) written on its sector 0 and also has a pvid attribute attached to the device attributes in the CuAt ODM. The **lspv** command lists all the physical volumes known to AIX. Here is a sample output from this command:

hdisk0	0001926922c706b2	rootvg
hdisk1	none	None
...		
hdisk10	none	None
hdisk11	00000000e7f5c88a	None
...		
hdisk48	none	None
hdisk49	00000000e7f5c88a	None
vpath0	00019269aa5bc858	None
vpath1	none	None
vpath2	none	None
vpath3	none	None
vpath4	none	None
vpath5	none	None
vpath6	none	None
vpath7	none	None
vpath8	none	None
vpath9	00019269aa5bbadd	vpathvg
vpath10	00019269aa5bc4dc	vpathvg
vpath11	00019269aa5bc670	vpathvg
vpath12	000192697f9fd2d3	vpathvg
vpath13	000192697f9fde04	vpathvg

In some cases, access to data is not lost, but failover protection might not be present. Failover protection can be lost in several ways:

- Losing a device path
- Creating a volume group from single-path SDD vpath devices
- A side effect of running the disk change method
- Running the **mksysb restore** command
- Manually deleting devices and running the configuration manager (cfgmgr)

The following sections provide more information about the ways that failover protection can be lost.

Losing a device path

Due to hardware errors, SDD might remove one or more nonfunctional paths from an SDD vpath device. The states of these nonfunctional paths are marked as `Dead`, `Invalid`, or `Close_Dead` by SDD. An SDD vpath device loses failover protection if it has only one functional path left. To determine if any of the SDD vpath devices have lost failover protection due to nonfunctional paths, use the **datapath query device** command to show the state of the paths to an SDD vpath device.

Creating a volume group from single-path SDD vpath devices

A volume group created using any single-path SDD vpath device does not have failover protection because there is no alternate path to the supported storage device LUN.

A side effect of running the disk change method

It is possible to modify attributes for an hdisk device by running the **chdev** command. The **chdev** command invokes the hdisk configuration method to make the requested change. In addition, the hdisk configuration method sets the `pvid` attribute for an hdisk if it determines that the hdisk has a `pvid` written on sector 0 of the LUN. This causes the SDD vpath device and one or more of its hdisks to have the same `pvid` attribute in the ODM. If the volume group containing the SDD vpath device is activated, the LVM uses the first device it finds in the ODM with the required `pvid` to activate the volume group.

As an example, if you issue the **lsvpcfg** command, the following output is displayed:

```
vpath0 (Avail pv vpathvg) 018FA067 = hdisk1 (Avail )
vpath1 (Avail ) 019FA067 = hdisk2 (Avail )
vpath2 (Avail ) 01AFA067 = hdisk3 (Avail )
vpath3 (Avail ) 01BFA067 = hdisk4 (Avail ) hdisk27 (Avail )
vpath4 (Avail ) 01CFA067 = hdisk5 (Avail ) hdisk28 (Avail )
vpath5 (Avail ) 01DFA067 = hdisk6 (Avail ) hdisk29 (Avail )
vpath6 (Avail ) 01EFA067 = hdisk7 (Avail ) hdisk30 (Avail )
vpath7 (Avail ) 01FFA067 = hdisk8 (Avail ) hdisk31 (Avail )
vpath8 (Avail ) 020FA067 = hdisk9 (Avail ) hdisk32 (Avail )
vpath9 (Avail pv vpathvg) 02BFA067 = hdisk20 (Avail ) hdisk44 (Avail )
vpath10 (Avail pv vpathvg) 02CFA067 = hdisk21 (Avail ) hdisk45 (Avail )
vpath11 (Avail pv vpathvg) 02DFA067 = hdisk22 (Avail ) hdisk46 (Avail )
vpath12 (Avail pv vpathvg) 02EFA067 = hdisk23 (Avail ) hdisk47 (Avail )
vpath13 (Avail pv vpathvg) 02FFA067 = hdisk24 (Avail ) hdisk48 (Avail )
```

The following example of a **chdev** command could also set the **pvid** attribute for an **hdisk**:

```
chdev -l hdisk46 -a pv=yes
```

For this example, the output of the **lsvpcfg** command would look similar to this:

```
vpath0 (Avail pv vpathvg) 018FA067 = hdisk1 (Avail )
vpath1 (Avail ) 019FA067 = hdisk2 (Avail )
vpath2 (Avail ) 01AFA067 = hdisk3 (Avail )
vpath3 (Avail ) 01BFA067 = hdisk4 (Avail ) hdisk27 (Avail )
vpath4 (Avail ) 01CFA067 = hdisk5 (Avail ) hdisk28 (Avail )
vpath5 (Avail ) 01DFA067 = hdisk6 (Avail ) hdisk29 (Avail )
vpath6 (Avail ) 01EFA067 = hdisk7 (Avail ) hdisk30 (Avail )
vpath7 (Avail ) 01FFA067 = hdisk8 (Avail ) hdisk31 (Avail )
vpath8 (Avail ) 020FA067 = hdisk9 (Avail ) hdisk32 (Avail )
vpath9 (Avail pv vpathvg) 02BFA067 = hdisk20 (Avail ) hdisk44 (Avail )
vpath10 (Avail pv vpathvg) 02CFA067 = hdisk21 (Avail ) hdisk45 (Avail )
vpath11 (Avail pv vpathvg) 02DFA067 = hdisk22 (Avail ) hdisk46 (Avail pv vpathvg)
vpath12 (Avail pv vpathvg) 02EFA067 = hdisk23 (Avail ) hdisk47 (Avail )
vpath13 (Avail pv vpathvg) 02FFA067 = hdisk24 (Avail ) hdisk48 (Avail )
```

The output of the **lsvpcfg** command shows that **vpath11** contains **hdisk22** and **hdisk46**. However, **hdisk46** is the one with the **pvid** attribute set. If you run the **lsvg -p vpathvg** command again, the output would look similar to this:

vpathvg:	PV_NAME	PV STATE	TOTAL PPs	FREE PPs	FREE DISTRIBUTION
	vpath10	active	29	4	00..00..00..00..04
	hdisk46	active	29	4	00..00..00..00..04
	vpath12	active	29	4	00..00..00..00..04
	vpath13	active	29	28	06..05..05..06..06

Notice that now device **vpath11** has been replaced by **hdisk46**. That is because **hdisk46** is one of the **hdisk** devices included in **vpath11** and it has a **pvid** attribute in the ODM. In this example, the LVM used **hdisk46** instead of **vpath11** when it activated volume group **vpathvg**. The volume group is now in a mixed mode of operation because it partially uses SDD **vpath** devices and partially uses **hdisk** devices. This is a problem that must be fixed because failover protection is effectively disabled for the **vpath11** physical volume of the **vpathvg** volume group.

Note: The way to fix this problem with the mixed volume group is to run the **dpvgfix vg-name** command after running the **chdev** command.

Manually deleting devices and running the configuration manager (cfgmgr)

In AIX 4.3 and 5.1, if there is more than one adapter and the storage device has been used as a physical volume of a volume group, running the **cfgmgr** command once might not configure all the hdisks. This means that if you manually delete SDD vpath devices and its hdisks (paths) and run configuration manager, you might lose the failover protection.

Assume that vpath3 is made up of hdisk4 and hdisk27 and that vpath3 is currently a physical volume. If the vpath3, hdisk4, and hdisk27 devices are all deleted by using the **rmdev** command and then **cfgmgr** is invoked at the command line, it is possible that only one path of the original vpath3 is configured by AIX. The following commands might produce this situation:

```
rmdev -d1 vpath3
rmdev -d1 hdisk4
rmdev -d1 hdisk27
cfgmgr
```

The **datapath query device** command might show now that only one path (either hdisk4 or hdisk27) is configured for vpath3.

To restore failover protection (that is, configure multiple paths for vpath3), complete the following steps:

1. Enter **cfgmgr** once for each installed SCSI or fibre-channel adapter, or enter **cfgmgr n** times, where *n* represents the number of paths per SDD device.
Tip: Running **cfgmgr n** times for n-path vpath configurations is not always required. It is only necessary to run **cfgmgr n** times for an n-path configuration if the supported storage device has been used as a physical volume of a volume group. This is because the AIX disk driver might configure only one set of hdisks from one adapter if pvid is present on a device.
2. Run **addpaths** to dynamically add the paths discovered by **cfgmgr** to SDD vpath devices.

The **addpaths** command allows you to dynamically add more paths to SDD vpath devices while they are in *Available* state. The **cfgmgr** command might need to be run *N* times when adding new LUNs.

This command opens a new path (or multiple paths) automatically if the SDD vpath device is in the *Open* state, and the original number of paths of the vpath is more than one. You can either use the Add Paths to Available Data Path Devices SMIT panel or run the **addpaths** command from the AIX command line. Go to “Dynamically adding paths to SDD vpath devices” on page 46 for more information about the **addpaths** command.

Using LVM with SDD

This section describes how to use the Logical Volume Manager (LVM) with SDD to import and export volume groups, recover from mixed volume groups, extend an existing volume group, and back up and restore files belonging to a volume group.

Importing volume groups with SDD

You can import a new volume group definition from a set of physical volumes with SDD vpath devices using the Volume Groups SMIT panel.

Note: To use this feature, you must either have root user authority or be a member of the system group.

Complete the following steps to import a volume group with SDD devices:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Import a Volume Group** and press **Enter**. The Import a Volume Group panel is displayed.
6. In the Import a Volume Group panel, complete the following tasks:
 - a. Enter the volume group that you want to import.
 - b. Enter the physical volume that you want to import.
 - c. Press **Enter** after making the changes.

You can press **F4** for a list of choices.

Exporting a volume group with SDD

You can export a volume group definition from the system with SDD vpath devices using the Volume Groups SMIT panel.

The **exportvg** command removes the definition of the volume group specified by the Volume Group parameter from the system. Because all system knowledge of the volume group and its contents are removed, an exported volume group is no longer accessible. The **exportvg** command does not modify any user data in the volume group.

A volume group is an unshared resource within the system; it should not be accessed by another system until it has been explicitly exported from its current system and imported on another. The primary use of the **exportvg** command, coupled with the **importvg** command, is to allow portable volumes to be exchanged between systems. Only a complete volume group can be exported, not individual physical volumes.

Using the **exportvg** command and the **importvg** command, you can also switch ownership of data on physical volumes shared between two systems.

Note: To use this feature, you must either have root user authority or be a member of the system group.

Complete the following steps to export a volume group with SDD devices:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Export a Volume Group** and press **Enter**. The Export a Volume Group panel is displayed.
6. Enter the volume group to export and press **Enter**.

You can use the **F4** key to select the volume group that you want to export.

Recovering from mixed volume groups

When an SDD volume group is not active (that is, varied off), and certain AIX system administrative operations cause a device reconfiguration, a `pvid` attribute will be created for the supported storage device `hdisks`. This will cause the SDD volume group to become a mixed volume group. The following command is an example of a command that does this:

```
chdev -l hdiskN -a queue_depth=30
```

Run the `dpovgfix` shell script to recover a mixed volume group. The syntax is `dpovgfix vg-name`. The script searches for an SDD `vpath` device that corresponds to each `hdisk` in the volume group and replaces the `hdisk` with the SDD `vpath` device. In order for the shell script to be run, all mounted file systems of this volume group have to be unmounted. After successful completion of the `dpovgfix` shell script, mount the file systems again.

Extending an existing SDD volume group

You can extend a volume group with SDD `vpath` devices using the Logical Volume Groups SMIT panel. The SDD `vpath` devices to be added to the volume group should be chosen from those that can provide failover protection. It is possible to add an SDD `vpath` device to an SDD volume group that has only a single path (`vpath0` on 73) and then add paths later by reconfiguring the supported storage device. With a single path, failover protection is not provided. (See “Dynamically adding paths to SDD `vpath` devices” on page 46 for information about adding paths to an SDD device.)

Complete the following steps to extend a volume group with SDD devices:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Group** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Add a Data Path Volume to a Volume Group** and press **Enter**.
6. Enter the volume group name and physical volume name and press **Enter**. You can also use the F4 key to list all the available SDD devices, and you can select the devices that you want to add to the volume group.

If you use a script file to extend an existing SDD volume group, you must modify your script file and replace the `extendvg` command with the `extendvg4vp` command.

Backing up all files belonging to an SDD volume group

You can back up all files belonging to a specified volume group with SDD `vpath` devices using the Volume Groups SMIT panel.

To back up a volume group with SDD devices, go to “Accessing the Backup a Volume Group with Data Path Devices SMIT panel” on page 85.

If you use a script file to back up all files belonging to a specified SDD volume group, you must modify your script file and replace the `savevg` command with the `savevg4vp` command.

Attention: Backing up files (running the **savevg4vp** command) will result in the loss of all material previously stored on the selected output medium. Data integrity of the archive might be compromised if a file is modified during system backup. Keep system activity at a minimum during the system backup procedure.

Restoring all files belonging to an SDD volume group

You can restore all files belonging to a specified volume group with SDD vpath devices using the Volume Groups SMIT panel.

To restore a volume group with SDD vpath devices, go to “Accessing the Remake a Volume Group with Data Path Devices SMIT panel” on page 86.

If you use a script file to restore all files belonging to a specified SDD volume group, you must modify your script file and replace the **restvg** command with the **restvg4vp** command.

SDD-specific SMIT panels

SDD supports several special SMIT panels. Some SMIT panels provide SDD-specific functions, while other SMIT panels provide AIX functions (but require SDD-specific commands). For example, the Add a Volume Group with Data Path Devices function uses the SDD **mkvg4vp** command, instead of the AIX **mkvg** command. Table 13 lists the SDD-specific SMIT panels and how you can use them.

Table 13. SDD-specific SMIT panels and how to proceed

SMIT panels	How to proceed using SMITTY – Go to:	Equivalent SDD command
Display Data Path Device Configuration	“Accessing the Display Data Path Device Configuration SMIT panel” on page 82	lsvpcfg
Display Data Path Device Status	“Accessing the Display Data Path Device Status SMIT panel” on page 83	datapath query device
Display Data Path Device Adapter Status	“Accessing the Display Data Path Device Adapter Status SMIT panel” on page 83	datapath query adapter
Define and Configure all Data Path Devices	“Accessing the Define and Configure All Data Path Devices SMIT panel” on page 84	cfallvpath
Add Paths to Available Data Path Devices	“Accessing the Add Paths to Available Data Path Devices SMIT panel” on page 84	addpaths
Configure a Defined Data Path Device	“Accessing the Configure a Defined Data Path Device SMIT panel” on page 84	mkdev
Remove a Data Path Device	“Accessing the Remove a Data Path Device SMIT panel” on page 84	rmdev
Add a Volume Group with Data Path Devices	“Accessing the Add a Volume Group with Data Path Devices SMIT panel” on page 84	mkvg4vp

Table 13. SDD-specific SMIT panels and how to proceed (continued)

Add a Data Path Volume to a Volume Group	"Accessing the Add a Data Path Volume to a Volume Group SMIT panel" on page 85	extendvg4vp
Remove a Physical Volume from a Volume Group	"Accessing the Remove a Physical Volume from a Volume Group SMIT panel" on page 85	exportvg <i>volume_group</i>
Back Up a Volume Group with Data Path Devices	"Accessing the Backup a Volume Group with Data Path Devices SMIT panel" on page 85	savevg4vp
Remake a Volume Group with Data Path Devices	"Accessing the Remake a Volume Group with Data Path Devices SMIT panel" on page 86	restvg

Accessing the Display Data Path Device Configuration SMIT panel

Complete the following steps to access the Display Data Path Device Configuration panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Configuration** and press **Enter**
5. The following example shows the Data Path Devices panel:

```

+-----+
|                                     |
|               Display Data Path Device Configuration               |
|                                     |
| Type or select values in entry fields.                             |
| Press Enter AFTER making all desired changes.                     |
|                                     |
| Select Query Option [Entry Fields]                                |
| Device Name/ Device Model all devices +                          |
|                                     [ ]                            |
|                                     |
+-----+

```

The Select Query Option has three options:

All devices

This option runs **lsvpcfg** and all the data path devices are displayed. No entry is required in the Device Name/Device Model field.

Device name

This option runs **lsvpcfg <device name>** and only the specified device is displayed. Enter a device name in the Device Name/Device Model field.

Device model

This option runs **lsvpcfg -d <device model>** and only devices with the specified device model are displayed. Enter a device model in the Device Name/Device Model field.

See “lsvpcfg” on page 87 for detailed information about the `lsvpcfg` command.

Accessing the Display Data Path Device Status SMIT panel

Complete the following steps to access the Display Data Path Device Status panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Status** and press **Enter**.
5. The following example shows the Data Path Devices Status panel:

```
+-----+
|                                     |
|               Display Data Path Device Status               |
| Type or select values in entry fields.                       |
| Press Enter AFTER making all desired changes.                |
|                                                                |
| Select Query Option                                         [Entry Fields] |
| Device Number/ Device Model                               all devices      +
|                                                                [ ]          |
+-----+
```

The Select Query Option has 3 options:

All devices

This option runs `datapath query device` and all the data path devices are displayed. No entry is required in the Device Name/Device Model field.

Device number

This option runs `datapath query device <device number>` and only the specified device is displayed. Enter a device number in the Device Name/Device Model field.

Device model

This option runs `datapath query device -d <device model>` and only devices with the specified device model are displayed. Enter a device model in the Device Name/Device Model field.

See “datapath query device” on page 378 for detailed information about the `datapath query device` command.

Accessing the Display Data Path Device Adapter Status SMIT panel

Complete the following steps to access the Display Data Path Device Adapter Status panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Adapter Status** and press **Enter**.

Accessing the Define and Configure All Data Path Devices SMIT panel

Complete the following steps to access the Define and Configure All Data Path Devices panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Define and Configure All Data Path Devices** and press **Enter**.

Accessing the Add Paths to Available Data Path Devices SMIT panel

Complete the following steps to access the Add Paths to Available Data Path Devices panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Add Paths to Available Data Path Devices** and press **Enter**.

Accessing the Configure a Defined Data Path Device SMIT panel

Complete the following steps to access the Configure a Defined Data Path Device panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Configure a Defined Data Path Device** and press **Enter**.

Accessing the Remove a Data Path Device SMIT panel

Complete the following steps to access the Remove a Data Path Device panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Remove a Data Path Device** and press **Enter**.

Accessing the Add a Volume Group with Data Path Devices SMIT panel

Complete the following steps to access the Add a volume group with data path devices panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.

3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Add Volume Group with Data Path Devices** and press **Enter**.

Note: Press **F4** while highlighting the **PHYSICAL VOLUME names** field to list all the available SDD vpaths.

Accessing the Add a Data Path Volume to a Volume Group SMIT panel

Complete the following steps to access the Add a Data Path Volume to a Volume Group panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical)** and press **Enter**. The System Storage Management (Physical & Logical) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Group** and press **Enter**. The Volume Group panel is displayed.
5. Select **Add a Data Path Volume to a Volume Group** and press **Enter**.
6. Enter the volume group name and physical volume name and press **Enter**. Alternately, you can use the **F4** key to list all the available SDD vpath devices and use the **F7** key to select the physical volumes that you want to add.

Accessing the Remove a Physical Volume from a Volume Group SMIT panel

Complete the following steps to access the Remove a Physical Volume from a Volume Group panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
3. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
4. Select **Set Characteristics of a Volume Group** and press **Enter**. The Set Characteristics of a Volume Group panel is displayed.
5. Select **Remove a Physical Volume from a Volume Group** and press **Enter**. The Remove a Physical Volume from a Volume Group panel is displayed.

Accessing the Backup a Volume Group with Data Path Devices SMIT panel

Complete the following steps to access the Back Up a Volume Group with Data Path Devices panel and to backup a volume group with SDD devices:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Back Up a Volume Group with Data Path Devices** and press **Enter**. The Back Up a Volume Group with Data Path Devices panel is displayed.

6. In the Back Up a Volume Group with Data Path Devices panel, complete the following steps:
 - a. Enter the Backup DEVICE or FILE name.
 - b. Enter the Volume Group to backup.
 - c. Press **Enter** after making all required changes.

Tip: You can also use the F4 key to list all the available SDD devices, and you can select the devices or files that you want to backup.

Attention: Backing up files (running the `savevg4vp` command) will result in the loss of all material previously stored on the selected output medium. Data integrity of the archive might be compromised if a file is modified during system backup. Keep system activity at a minimum during the system backup procedure.

Accessing the Remake a Volume Group with Data Path Devices SMIT panel

Complete the following steps to access the Remake a Volume Group with Data Path Devices panel and restore a volume group with SDD devices:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Remake a Volume Group with Data Path Devices** and press **Enter**. The Remake a Volume Group with Data Path Devices panel is displayed.
6. Enter the Restore DEVICE or FILE name that you want to restore, and press **Enter**. You can also press **F4** to list all the available SDD devices, and you can select the devices or files that you want to restore.

SDD utility programs

The following SDD utility programs are available:

addpaths

You can use the **addpaths** command to dynamically add more paths to the SDD devices when they are in the *Available* state. In addition, this command allows you to add paths to the SDD vpath devices (which are then opened) belonging to active volume groups.

This command will open a new path (or multiple paths) automatically if the SDD vpath device is in *Open* state. You can either use the Add Paths to Available Data Path Devices SMIT panel or run the **addpaths** command from the AIX command line.

The syntax for this command is:

```
▶▶—addpaths—▶▶
```

For more information about this command, go to “Dynamically adding paths to SDD vpath devices” on page 46.

hd2vp and vp2hd

The SDD provides two conversion scripts, **hd2vp** and **vp2hd**. The **hd2vp** script converts a volume group from supported storage device hdisks to the SDD vpath devices, and the **vp2hd** script converts a volume group from the SDD vpath devices to supported storage device hdisks. Use the **vp2hd** program when you want to configure your applications back to original supported storage device hdisks, or when you want to remove the SDD from your AIX host system.

The syntax for these conversion scripts is as follows:

```
▶▶—hd2vp—vname—————▶▶
```

```
▶▶—vp2hd—vname—————▶▶
```

vname

Specifies the volume group name to be converted.

dpovgfix

You can use the **dpovgfix** script tool to recover mixed volume groups.

Performing AIX system management operations on adapters and hdisk devices can cause original supported storage device hdisks to be contained within an SDD volume group. This is known as a mixed volume group. Mixed volume groups happen when an SDD volume group is not active (varied off), and certain AIX commands to the hdisk put the **pvid** attribute of hdisk back into the ODM database. The following is an example of a command that does this:

```
chdev -l hdiskW -a queue_depth=30
```

If this disk is an active hdisk of an SDD vpath device that belongs to an SDD volume group, and you run the **varyonvg** command to activate this SDD volume group, LVM might pick up the hdisk device instead of the SDD vpath device. The result is that an SDD volume group partially uses the SDD vpath devices, and partially uses supported storage device hdisk devices. This causes the volume group to lose path-failover capability for that physical volume. The **dpovgfix** script tool fixes this problem.

The syntax for this command is:

```
▶▶—dpovgfix—vname—————▶▶
```

vname

Specifies the volume group name of the mixed volume group to be recovered.

lsvpcfg

You can use the **lsvpcfg** script tool to display the configuration state of SDD devices. This displays the configuration state for all SDD devices. The **lsvpcfg** command can be issued in three ways.

1. The command can be issued without parameters. The syntax for this command is:

lsvpcfg

See “Verifying the SDD configuration” on page 45 for an example of the output and what it means.

2. The command can also be issued using the SDD vpath device name as a parameter. The syntax for this command is:

lsvpcfg vpathW₀ vpathW₁ vpathW₂

You will see output similar to this:

```
vpath10 (Avail pv ) 13916392 = hdisk95 (Avail ) hdisk179 (Avail )
vpath20 (Avail ) 02816392 = hdisk23 (Avail ) hdisk106 (Avail )
vpath30 (Avail ) 10516392 = hdisk33 (Avail ) hdisk116 (Avail )
```

See “Verifying the SDD configuration” on page 45 for an explanation of the output.

3. The command can also be issued using the device model as a parameter. The option to specify a device model cannot be used when you specify an SDD vpath device. The syntax for this command is:

lsvpcfg device model

The following are examples of valid device models:

2105 All 2105 models (ESS).

2105F All 2105 F models (ESS).

2105800

All 2105 800 models (ESS).

2145 All 2145 models (SAN Volume Controller).

2107 All DS8000 models.

1750 All DS6000 models.

mkvg4vp

You can use the **mkvg4vp** command to create an SDD volume group. For more information about this command, go to “Configuring volume groups for failover protection” on page 74. For information about the flags and parameters for this command, go to:

publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/cmds/aixcmds3/mkvg.htm.

The syntax for this command is:

```

▶▶ mkvg4vp [ -S* ] [ -d MaxPVs ] [ -B ] [ -G ] [ -f ] [ -q ** ]
▶ [ -C | -c [-x] ] [ -i ] [ -s PPSize ] [ -n ]
▶ [ -m MaxPVsize | -t factor ] [ -V MajorNumber ] [ -L LTGsize*** ]

```

- * for AIX 5.3 and later only
- ** for AIX 5.2 and later only
- *** for AIX 5.1 and later only

extendvg4vp

You can use the **extendvg4vp** command to extend an existing SDD volume group. For more information about this command, go to “Extending an existing SDD volume group” on page 80. For information about the flag and parameters for this command, go to:

publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/cmds/aixcmds2/extendvg.htm

The syntax for this command is:

excludesddcfg

You can use the **excludesddcfg** command to exclude supported storage device (hdisk) from the SDD vpath configuration. You must run this command before the SDD vpath devices are configured. The **excludesddcfg** command saves the serial number for the logical device (hdisk) to an exclude file (/etc/vpexclude). During the SDD configuration, the SDD configure methods read all serial numbers that are listed in this file and exclude these devices from the SDD configuration.

The syntax for this command is:

- l Specifies the logical number of the supported storage device (hdiskN). This is not the SDD device name.
- d When this optional flag is set, the **excludesddcfg** command removes all the exclusions by deleting all the existing contents from the exclude file.
- dl When this optional flag is set, the **excludesddcfg** command allows users to remove the exclusion of a particular storage device.

device name

Specifies the supported storage device (hdiskN).

Example:

```
# excludesddcfg -l hdisk11
```

```
hdisk11 SERIAL_NUMBER = 7C0FCA30
```

Success: Device with this serial number is now excluded from the SDD configuration. To undo this exclusion, run `excludesddcfg -dl hdisk#`.

```
# excludesddcfg -dl hdisk11
hdisk11 SERIAL NUMBER = 7C0FCA30 TYPE = 2105
```

Success: SERIAL_NUMBER 7C0FCA30 is removed from /etc/vpexclude file.
To configure previously excluded device(s), run cfallvpath.

Note:

1. Do not use the **excludesddcfg** command to exclude a device if you want the device to be configured by SDD.
2. If the supported storage device LUN has multiple configurations on a server, use the **excludesddcfg** command on only one of the logical names of that LUN.
3. Do not use the **excludesddcfg** command multiple times on the same logical device. Using the **excludesddcfg** command multiple times on the same logical device results in duplicate entries in the /etc/vpexclude file, so that the system administrator has to administer the file and its content.
4. Issue the **excludesddcfg** command with the -d flag to delete all existing contents from the exclude file. If you want to remove only one device from the /etc/vpexclude file, issue the **excludesddcfg** command with the -dl flag and specify the logical device name for which you want to remove the exclusion. For detailed instructions on the proper procedure, see “Replacing manually excluded devices in the SDD configuration” on page 52.

lquerypr

See “Persistent reserve command tool.”

sddgetdata

See Appendix A, “SDD, SDDPCM, and SDDDSM data collection for problem analysis,” on page 393, which describes the use of sddgetdata to collect information for problem determination.

Persistent reserve command tool

SDD offers the **lquerypr** tool to provide a set of persistent reserve functions. This tool supports the following persistent reserve service actions:

- Read persistent reservation key
- Release persistent reserve
- Preempt-abort persistent reserve
- Clear persistent reserve and registrations

Note: Caution must be taken with the command, especially when implementing preempt-abort or clear persistent reserve service action. With preempt-abort service action not only the current persistent reserve key is preempted; it also aborts tasks on the LUN that originated from the initiators that are registered with the preempted key. With clear service action, both persistent reservation and reservation key registrations are cleared from the device or LUN.

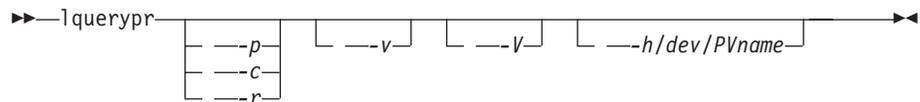
The following information describes in detail the syntax and examples of the **lquerypr** command.

lquerypr command

Purpose

To query and implement certain SCSI-3 persistent reserve commands on a device.

Syntax



Description

The **lquerypr** command implements certain SCSI-3 persistent reservation commands on a device. The device can be either hdisk or SDD vpath devices. This command supports persistent reserve service actions, such as read reservation key, release persistent reservation, preempt-abort persistent reservation, and clear persistent reservation.

Note: This command can only be used when the device is not already opened.

Flags

- p** If the persistent reservation key on the device is different from the current host reservation key, it preempts the persistent reservation key on the device.
- c** If there is a persistent reservation key on the device, it removes any persistent reservation and clears all reservation key registration on the device.
- r** Removes the persistent reservation key on the device made by this host.
- v** Displays the persistent reservation key if it exists on the device.
- V** Verbose mode. Prints detailed message.

Return code

If the command is issued without options of **-p**, **-r**, or **-c**, the command will return 0 under two circumstances.

1. There is no persistent reservation key on the device.
2. The device is reserved by the current host.

If the persistent reservation key is different from the host reservation key, the command will return 1. If the command fails, it returns 2. If the device is already opened on a current host, the command returns 3.

Example

1. To query the persistent reservation on a device, enter **lquerypr -h/dev/vpath30**.

This command queries the persistent reservation on the device without displaying. If there is a persistent reserve on a disk, it returns 0 if the device is reserved by the current host. It returns 1 if the device is reserved by another host.

2. To query and display the persistent reservation on a device, enter **lquerypr -vh/dev/vpath30**.

Same as Example 1. In addition, it displays the persistent reservation key.

3. To release the persistent reservation if the device is reserved by the current host, enter **lquerypr -rh/dev/vpath30**.

This command releases the persistent reserve if the device is reserved by the current host. It returns 0 if the command succeeds or the device is not reserved. It returns 2 if the command fails.

4. To reset any persistent reserve and clear all reservation key registrations, enter **lquerypr -ch/dev/vpath30**.

This command resets any persistent reserve and clears all reservation key registrations on a device. It returns 0 if the command succeeds, or 2 if the command fails.

5. To remove the persistent reservation if the device is reserved by another host, enter **lquerypr -ph/dev/vpath30**.

This command removes an existing registration and persistent reserve from another host. It returns 0 if the command succeeds or if the device is not persistent reserved. It returns 2 if the command fails.

Using supported storage devices directly

When you configure your host adapters, the AIX disk driver creates hdisks for the supported storage device LUNs. When you configure SDD, it creates SDD vpath devices for supported storage device LUNs. Therefore, with SDD installed on a system, an application now has two ways to access the supported storage devices.

To use the SDD load-balancing and failover features to access supported storage devices, your application must use SDD vpath devices only rather than hdisks.

Applications can access SDD vpath devices either by using raw devices or logical devices. If the application is accessing SDD vpath devices using logical devices, a volume group must be created with SDD vpath devices.

If your application used hdisk special files directly before installing SDD, convert the application to use SDD vpath device special files. After installing SDD, complete the following steps:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Configuration** and press **Enter**.
5. To display all SDD vpath devices with their attached multiple paths (hdisks), select all SDD vpath devices for **Select Query Option**, leave Device Name/Device Model blank, and press **Enter**.
6. Search the list of hdisks to locate the hdisks that your application is using.
7. Replace each hdisk with its corresponding SDD vpath device.

Note: Depending upon your application, the manner in which you replace these files is different. If this is a new application, use the SDD vpath device rather than hdisk to use the SDD load-balancing and failover features.

Note: Alternately, you can enter **lsvpcfg** from the command-line interface rather than using SMIT. This displays all configured SDD vpath devices and their underlying paths (hdisks).

Using supported storage devices through AIX LVM

If your application accesses supported storage devices through LVM, determine that the physical volumes of the volume group that the application is accessing are SDD-supported storage devices. Then complete the following steps to convert the volume group from the original supported storage device hdisks to the SDD vpath devices:

1. Determine the file systems or logical volumes that your application accesses.
2. Enter `smitty` from your desktop window. SMIT is displayed.
3. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
4. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
5. Select **Logical Volume** and press **Enter**. The Logical Volume panel is displayed.
6. Select **List All Logical Volumes by Volume Group** to determine the logical volumes that belong to this volume group and their logical volume mount points.
7. Press **Enter**. The logical volumes are listed by volume group.
To determine the file systems, complete the following steps:
 - a. Enter `smitty` from your desktop window. SMIT is displayed.
 - b. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
 - c. Select **File Systems** and press **Enter**. The File Systems panel is displayed.
 - d. Select **List All File Systems** to locate all file systems that have the same mount points as the logical volumes and press **Enter**. The file systems are listed.
 - e. Note the file system name of that volume group and the file system mount point, if it is mounted.
 - f. Unmount these file systems.
8. Enter the following command to convert the volume group from the supported storage device hdisks to SDD multipath vpath devices:

```
hd2vp vgroupname
```
9. When the conversion is complete, mount all file systems that you previously unmounted.

When the conversion is complete, your application now accesses supported storage device physical LUNs through SDD vpath devices. This provides load-balancing and failover protection for your application.

Migrating a non-SDD volume group to a supported storage device SDD multipath volume group in concurrent mode

Before you migrate your non-SDD volume group to an SDD volume group, make sure that you have completed the following tasks:

- The SDD for the AIX host system is installed and configured. See “Verifying the currently installed version of SDD for SDD 1.3.3.11 (or earlier)” on page 35 or “Verifying the currently installed version of SDD for SDD 1.4.0.0 (or later)” on page 36.

- The supported storage devices to which you want to migrate have multiple paths configured per LUN. To check the state of your SDD configuration, use the System Management Interface Tool (SMIT) or issue the **lsvpcfg** command from the command line. To use SMIT:
 - Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
 - Select **Devices** and press **Enter**. The Devices panel is displayed.
 - Select **Data Path Device** and press **Enter**. The Data Path Device panel is displayed.
 - Select **Display Data Path Device Configuration** and press **Enter**. A list of the SDD vpath devices and whether there are multiple paths configured for the devices is displayed.
- Ensure that the SDD vpath devices that you are going to migrate to do not belong to any other volume group, and that the corresponding physical device (supported storage device LUN) does not have a pvid written on it. Enter the **lsvpcfg** command output to check the SDD vpath devices that you are going to use for migration. Make sure that there is no pv displayed for this SDD vpath device and its paths (hdisks). If a LUN has never belonged to any volume group, there is no pvid written on it. In case there is a pvid written on the LUN and the LUN does not belong to any volume group, you need to clear the pvid from the LUN before using it to migrate a volume group. The commands to clear the pvid are:

```
chdev -l hdiskW -a pv=clear
chdev -l vpathW -a pv=clear
```

Attention: Exercise care when clearing a pvid from a device with this command. Issuing this command to a device that *does* belong to an existing volume group can cause system failures.

You should complete the following steps to migrate a non-SDD volume group to a multipath SDD volume group in concurrent mode:

1. Add new SDD vpath devices to an existing non-SDD volume group:
 - a. Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
 - b. Select **System Storage Management (Physical & Logical)** and press **Enter**. The System Storage Management (Physical & Logical) panel is displayed.
 - c. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
 - d. Select **Volume Group** and press **Enter**. The Volume Group panel is displayed.
 - e. Select **Add a Data Path Volume to a Volume Group** and press **Enter**.
 - f. Enter the volume group name and physical volume name and press **Enter**. Alternately, you can use the F4 key to list all the available SDD vpath devices and use the F7 key to select the physical volumes that you want to add.
2. Enter the **smitty mklvcopy** command to mirror logical volumes from the original volume to an SDD supported storage device volume. Use the new SDD vpath devices for copying all logical volumes. Do not forget to include JFS log volumes.

Note: The command **smitty mklvcopy** copies one logical volume at a time. A fast-path command to mirror *all* the logical volumes on a volume group is **mirrorvg**.

3. Synchronize logical volumes (LVs) or force synchronization. Enter the **smitty syncvg** command to synchronize all the volumes:
There are two options on the smitty panel:
 - Synchronize by Logical Volume
 - Synchronize by Physical Volume
 The fast way to synchronize logical volumes is to select the **Synchronize by Physical Volume** option.
4. Remove the mirror and delete the original LVs. Enter the **smitty rmlvcopy** command to remove the original copy of the logical volumes from all original non-SDD physical volumes.
5. Enter the **smitty reducevg** command to remove the original non-SDD vpath devices from the volume group. The Remove a Physical Volume panel is displayed. Remove all non-SDD devices.

Note: A non-SDD volume group refers to a volume group that consists of non-supported storage devices or supported storage hdisk devices.

Detailed instructions for migrating a non-SDD volume group to a supported storage device SDD multipath volume group in concurrent mode

This procedure shows how to migrate an existing AIX volume group to use SDD vpath devices that have multipath capability. You do not take the volume group out of service. The example shown starts with a volume group, `vg1`, made up of one supported storage device, `hdisk13`.

To perform the migration, you must have SDD vpath devices available that are greater than or equal to the size of each of the hdisks making up the volume group. In this example, the volume group is migrated to an SDD device, `vpath12`, with two paths, `hdisk14` and `hdisk30`.

1. Add the SDD vpath device to the volume group as an Available volume:
 - a. Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
 - b. Select **System Storage Management (Physical & Logical)** and press **Enter**. The System Storage Management (Physical & Logical) panel is displayed.
 - c. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
 - d. Select **Volume Group** and press **Enter**. The Volume Group panel is displayed.
 - e. Select **Add a Data Path Volume to a Volume Group** and press **Enter**.
 - f. Enter **vg1** in the **Volume Group Name** field and enter **vpath12** in the **Physical Volume Name** field. Press **Enter**.
You can also use the **extendvg4vp -f vg1 vpath12** command.
2. Mirror logical volumes from the original volume to the new SDD vpath device volume:
 - a. Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
 - b. Select **System Storage Management (Physical & Logical)** and press **Enter**. The System Storage Management (Physical & Logical) panel is displayed.
 - c. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.

- d. Select **Volume Group** and press **Enter**. The Volume Group panel is displayed.
 - e. Select **Mirror a Volume Group** and press **Enter**. The Mirror a Volume Group panel is displayed.
 - f. Enter a volume group name and a physical volume name. Press **Enter**. You can also enter the **mirrorvg vg1 vpath12** command.
3. Synchronize the logical volumes in the volume group:
 - a. Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
 - b. Select **System Storage Management (Physical & Logical)** and press **Enter**. The System Storage Management (Physical & Logical) panel is displayed.
 - c. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
 - d. Select **Volume Group** and press **Enter**. The Volume Group panel is displayed.
 - e. Select **Synchronize LVM Mirrors** and press **Enter**. The Synchronize LVM Mirrors panel is displayed.
 - f. Select **Synchronize by Physical Volume**.

You can also enter the **syncvg -p hdisk13 vpath12** command.

4. Delete copies of all logical volumes from the original physical volume:
 - a. Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
 - b. Select **Logical Volumes** and press **Enter**. The Logical Volumes panel is displayed.
 - c. Select **Set Characteristic of a Logical Volume** and press **Enter**. The Set Characteristic of a Logical Volume panel is displayed.
 - d. Select **Remove Copy from a Logical Volume** and press **Enter**. The Remove Copy from a Logical Volume panel is displayed.

You can also enter the command:

```
rm1vcopy loglv01 1 hdisk13
rm1vcopy lv01 1 hdisk13
```

5. Remove the old physical volume from the volume group:
 - a. Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
 - b. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
 - c. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
 - d. Select **Set Characteristics of a Volume Group** and press **Enter**. The Set Characteristics of a Volume Group panel is displayed.
 - e. Select **Remove a Physical Volume from a Volume Group** and press **Enter**. The Remove a Physical Volume from a Volume Group panel is displayed.

You can also enter the **reducevg vg1 hdisk13** command.

Using the trace function

SDD supports AIX trace functions. The SDD trace ID is 2F8. Trace ID 2F8 traces routine entry, exit, and error paths of the algorithm. To use it, manually turn on the trace function before the program starts to run, then turn off the trace function either after the program stops, or any time you need to read the trace report.

By default, SDD disables the AIX trace functions. To enable the AIX trace functions for SDD, follow these steps:

1. Enter **pathctest -d** . (for example, **pathctest -d 0**)
2. Enter **777**.
3. Enter **20** to open the device.
4. Enter **3** (as option NO_DELAY).
5. Enter **90** (enable or disable the AIX trace). Follow the prompt: enter **1** to enable.

Then you can start the trace function.

To start the trace function, enter:

```
trace -a -j 2F8
```

To stop the trace function, enter:

```
trcstop
```

To read the report, enter:

```
trcrpt | pg
```

To save the trace data to a file, enter:

```
trcrpt > filename
```

Note: To perform the AIX trace function, you must have the bos.sysmgt.trace installation package installed on your system.

Chapter 3. Using SDDPCM on an AIX host system

SDDPCM is a loadable path control module for supported storage devices to supply path management functions and error recovery algorithms. When the supported storage devices are configured as Multipath I/O (MPIO) devices, SDDPCM is loaded as part of the AIX MPIO FCP (Fibre Channel Protocol) or AIX MPIO SAS (serial-attached SCSI) device driver during the configuration. The AIX MPIO-capable device driver with the supported storage devices SDDPCM module enhances the data availability and I/O load balancing.

This chapter provides a general view of the SDDPCM path control module, including where it resides on the I/O stack in the operating system and the features and functions that it supports. This chapter also provides procedures for the following management tasks:

- Install SDDPCM
- Configure SDDPCM MPIO-capable devices
- Uninstall SDDPCM
- Migrate disk storage MPIO-capable devices from the AIX default PCM to SDDPCM
- Migrate disk storage MPIO-capable devices from SDDPCM to the AIX default PCM or SDD
- Installing SDDPCM from AIX NIM SPOT to client SAN boot disk or internal boot disk

Figure 3 on page 100 shows the position of SDDPCM in the protocol stack. I/O operations are sent to the AIX disk driver. The SDDPCM path selection routine is invoked to select an appropriate path for each I/O operation.

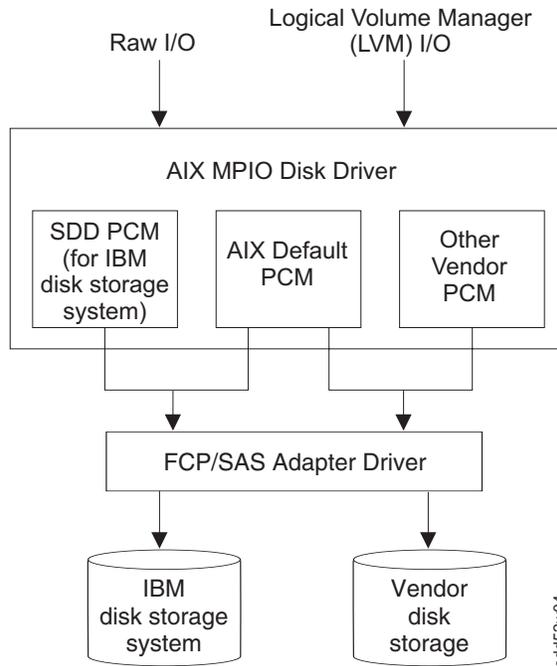


Figure 3. SDDPCM in the protocol stack

For detailed information about MPIO support on AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1, contact the IBM Support team.

AIX MPIO-capable device drivers automatically discover, configure and make available every storage device path. SDDPCM manages the paths to provide:

- High availability and load balancing of storage I/O
- Automatic path-failover protection
- Concurrent download of supported storage devices licensed machine code
- Prevention of a single-point-failure

For updated and additional information that is not included in this chapter, see the Readme file on the CD-ROM or visit the SDD website:

www-947.ibm.com/support/entry/portal/product/system_storage/storage_software/other_software_products/system_storage_multipath_subsystem_device_driver?productContext=-1564959213

SDD and SDDPCM are exclusive software packages on a server. You cannot install both software packages on a server for supported storage devices. When supported storage devices are configured as non-MPIO-capable devices (that is, multiple logical device instances are created for a physical LUN), you should install SDD to get multipath support.

You must install SDDPCM in order to configure supported storage devices into MPIO-capable-devices (where only one logical device instance is created for a physical LUN). Before you install SDDPCM, make sure that you meet all the required hardware and software requirements. See “Verifying the hardware and software requirements” on page 103 and the “Preparing for SDDPCM installation for supported storage devices” on page 106.

Note: SDDPCM does not support SCSI storage devices.

With SDD 1.6.0.0 (or later), SDDPCM and SDD cannot coexist on a AIX server. If a server connects to any supported storage devices, all devices must be configured either as non-MPIO-capable devices or as MPIO-capable devices.

Supported SDDPCM features

The following SDDPCM features are supported in this release:

- 32-bit and 64-bit kernels
- Four types of reserve policies:
 - No reserve policy
 - Exclusive host access single path policy
 - Persistent reserve exclusive host policy
 - Persistent reserve shared host access policy
- Four path-selection algorithms:
 - Failover
 - Round robin
 - Load balancing
 - Load balancing port

This algorithm is supported only with DS8000, DS6000, ESS, and SAN Volume Controller device models.
- Automatic failed paths reclamation by healthchecker
- Failback error-recovery algorithm
- Fibre-channel dynamic device tracking
- Support for all DS4000, DS5000, DS3950, DS6000, DS8000, ESS FCP, and SAN Volume Controller devices
- Support SAN boot device on MPIO supported storage devices

Note:

1. SDDPCM supports ESS devices as SAN boot devices, beginning with AIX52 TL06 and AIX53 TL02.
 2. SDDPCM supports DS8000, DS6000, and SAN Volume Controller devices as SAN boot devices, beginning with AIX52 TL07 and AIX53 TL03.
 3. SDDPCM supports DS4000, DS5000, and DS3950 devices as SAN boot devices, beginning with AIX53 TL8 and AIX61 TL2
 4. SDDPCM supports RSSM devices as SAN boot devices, beginning with AIX61 TL3.
- Support for external supported MPIO storage devices as the primary or secondary dump device
 - Support storage devices multipath devices as system paging space
 - Support SDDPCM server daemon enhanced path health check function
 - Support for a maximum of 65536 LUNs for SDDPCM version 2.6.5.0 or later on an AIX system with sufficient resources to handle higher number of LUNs. For earlier versions of SDDPCM, a maximum of 1200 LUNs are supported.

Note: For more information about devices that SDDPCM supports, see “Maximum number of devices that SDDPCM supports” on page 116.

- Dynamically adding paths or adapters
- Dynamically removing paths or adapters
- Dynamically changing the device path selection algorithm

- Dynamically changing the device `hc_interval`
- Dynamically enabling or disabling the device health check function
- Dynamically changing the device `hc_mode`
- Web-based System Manager (WebSM) for MPIO supported storage devices (See www-03.ibm.com/systems/power/software/aix/index.html for more information about WebSM.)
- Reserve last path of a device in OPEN mode
- Support the `essutil` Product Engineering tool in the SDDPCM `pcmpath` command line program
- Support for Live Partition Mobility (LPM) environment in IBM Power servers. For information about LPM, see <http://www.redbooks.ibm.com/abstracts/sg247460.html>.
- Support HACMP with Enhanced Concurrent Mode volume group in concurrent resource groups and nonconcurrent resource groups

Note: This support does not include RSSM devices as HACMP is not supported on IBM JS-series blades.

- Support GPFS in AIX 5.2 TL06 (or later), 5.3 TL02 (or later), and AIX 6.1.
- Support Virtual I/O server with AIX 5.3 or later and AIX 6.1.

Note: This support does not include DS4000, DS5000, and DS3950 storage devices.

- Support Tivoli Storage Productivity Center for Replication Metro Mirror Failover/Failback replication for Open HyperSwap for System Storage DS8000.

Note: This support is only available on AIX 5.3 TL11 (or later), and AIX 6.1 TL04 (or later).

- Support for Non-disruptive Vdisk Movement (NDVM) feature of SVC. For information about SAN Volume Controller, see the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

Note: This support is only available on SDDPCM 2.6.4.0 (or later) on AIX.

- Support for Workload partitions (WPAR) and Versioned WPAR (vWPAR) for AIX. For information about WPAR and vWPAR, see publib.boulder.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.wpar/wpar-kickoff.htm.

Note: This support is only available on SDDPCM 2.6.5.0 (or later) on AIX.

- Support for stretched cluster environment of SVC. For information about stretched clusters, see pic.dhe.ibm.com/infocenter/svc/ic/topic/com.ibm.storage.svc.console.720.doc/svc_stretchedclusteroverview.html.

Note: This support is only available on SDDPCM 2.6.5.0 (or later) on AIX.

Unsupported SDDPCM features

The following SDDPCM features are not currently supported.

- HACMP with persistent reservation policies
- Multiple hosts sharing supported storage MPIO devices without supported clustering software installed
- HACMP with non-Enhanced Concurrent Mode volume group
- Mirror Write Consistency with volume group in a concurrent resource group

- Virtual I/O server is not supported with DS4000, DS5000, and DS3950 storage subsystems
- Open HyperSwap for System Storage DS8000 is not supported within a clustering environment.
- Open HyperSwap device is not supported with a SAN boot volume group.

Verifying the hardware and software requirements

You must install the following hardware and software components to ensure that SDDPCM installs and operates successfully.

Hardware

The following hardware components are needed:

- Supported storage devices (FCP and SAS devices only)
- One or more switches, if the supported storage devices is not direct-attached
- Host system
- Fibre-channel adapters or serial-attached SCSI adapters, and cables

Software

The following software components are needed:

- AIX 5.2 TL10 (or later), AIX 5.3 TL08 (or later), or AIX 6.1 TL02 (or later) operating system, with all the latest PTFs. Refer to the Readme file of the SDDPCM level that you plan to install for the required AIX TL for that level.
- If your attached storage is SAN Volume Controller version 4.2.1.6 or above and you require SAN Volume Controller APAR IC55826, you must install SDDPCM 2.2.0.3 or above with the required AIX TL and APAR. Refer to the Readme file of the SDDPCM level that you plan to install for the required AIX TL and APARs.
- Fibre-channel device drivers or serial-attached SCSI drivers
- One of the following installation packages:
 - devices.sddpcm.52.rte (version 2.5.1.0 or later)
 - devices.sddpcm.53.rte (version 2.5.1.0 or later)
 - devices.sddpcm.61.rte (version 2.5.1.0 or later)
- Supported storage devices:
 - devices.fcp.disk.ibm.mpio.rte (version 1.0.0.21 or later versions of 1.x.x.x) host attachment package for SDDPCM (version 2.x.x.x or prior versions)
 - devices.fcp.disk.ibm.mpio.rte (version 2.0.0.1 or later versions) host attachment package for SDDPCM (version 3.0.0.0 or later versions)
 - devices.sas.disk.ibm.mpio.rte (version 1.0.0.0 or later versions) host attachment package for SDDPCM

Unsupported environments

SDDPCM does not support:

- ESS SCSI devices
- A host system with both a SCSI and fibre-channel connection to a shared ESS logical unit number (LUN)
- Single-path mode during code distribution and activation of LMC nor during any supported storage devices concurrent maintenance that impacts the path attachment, such as a supported storage device host-bay-adapter replacement.
- AIX5.2 with DS4000, DS5000, and DS3950 storage subsystems.

- AIX5.2 with Open HyperSwap.

Host system requirements

To successfully install SDDPCM for supported storage devices, you must have the required AIX TL installed on your host system along with the AIX required fixes, APARs, and microcode updates that are identified on the following website:

www-933.ibm.com/support/fixcentral/

Starting with SDDPCM 2.2.0.0, SDDPCM requires AIX 5.2 TL10 (or later), AIX 5.3 TL06 (or later), or AIX 6.1, with all the latest PTFs. Always refer to the Readme file for the required AIX TL for the SDDPCM level that you plan to install.

Supported storage device requirements

To successfully install SDDPCM, ensure that the `devices.fcp.disk.ibm.mpio.rte`, or `devices.sas.disk.ibm.mpio.rte` package, is installed on the server. Refer to “Installing and upgrading the AIX SDDPCM host attachment” on page 109 for more information.

Fibre requirements

Note: There is no fibre requirement for RSSM devices that are connected through SAS.

You must check for and download the latest fibre-channel device driver APARs, maintenance-level fixes, and microcode updates from the following website:

www-933.ibm.com/support/fixcentral/

If your host has only one fibre-channel adapter, it requires you to connect through a switch to multiple supported storage device ports. You should have at least two fibre-channel adapters to prevent data loss due to adapter hardware failure or software failure.

For information about the fibre-channel adapters that can be used on your AIX host system, go to the following website:

www.ibm.com/servers/storage/support

To use the SDDPCM fibre-channel support, ensure that your host system meets the following requirements:

- The AIX host system is an IBM RS/6000 or IBM System p with AIX 5.2 TL10 (or later), AIX 5.3 TL06 (or later), or AIX 6.1.
- The AIX host system has the fibre-channel device drivers installed along with all latest APARs.
- The host system can be a single processor or a multiprocessor system, such as SMP.
- A fiber-optic cable connects each fibre-channel adapter to a supported storage system port.
- If you need the SDDPCM I/O load-balancing and failover features, ensure that a minimum of two paths to a device are attached.

Preparing for SDDPCM installation

The SDDPCM installation package installs a number of major files on your AIX system. The major files that are part of the SDDPCM installation package are:

File name

Description

sddpcmrtl

A dynamically loaded module added to the device configuration methods to extend the supported storage devices device configuration methods to facilitate the configuration operations of the PCM KE

sddpcmke

A dynamically-loaded module added to the AIX 5L kernel that provides path management functions for supported storage devices

sdduserke

A dynamically-loaded module added to the AIX 5L kernel that provides the API to sddpcmke

pcmpath

SDDPCM command line tool

pcmsrv

Daemon for enhanced path healthcheck, and First Time Data Capture

sample_pcmsrv.conf

The sample SDDPCM server daemon configuration file

fcppcmmap

Collects supported storage devices fibre-channel device information through SCSI commands

pcmquerypr

SDDPCM persistent reserve command tool

pcmgenprkey

SDDPCM persistent reserve command tool to generate persistent reserve key

relbootrsv

Release SCSI-2 reserve on boot devices or on active nonboot volume groups

Note: Beginning with SDDPCM version 2.5.1.0, the relbootrsv program also can be used to release SCSI-2 reserve on an MPIO hdisk.

relSDDPCMbootrsv

Automatically executed upon reboot to release leftover SCSI-2 reserve on SDDPCM boot device, beginning with SDDPCM version 2.2.0.1.

sddpcmgetdata

Script to collect SDDPCM information, trace log files, and system error logs into an sddpcmdata_*host_date_time*.tar file for problem determination

sddpcm_get_config

Display configuration information for DS4000, DS5000, and DS3950 storage and MPIO-based devices.

lspcmcfg

Display SDDPCM MPIO device/path configuration/state information.

- AE** Daemon for communication with Tivoli Storage Productivity Center for Replication to support Open HyperSwap.

Note: This file was added to the SDDPCM installation package beginning with SDDPCM 3.0.0.0.

Preparing for SDDPCM installation for supported storage devices

Before you install SDDPCM, you must complete the following tasks:

- Connect the supported storage device to your host system and the required fibre-channel adapters or SAS adapters that are attached.
- Configure the supported storage device for single-port or multiple-port access for each LUN. To use the load-balancing and failover features, SDDPCM requires a minimum of two independent paths that share the same logical unit. **With a single path, failover protection is not available.**

For more information about how to configure your supported storage device, see the Planning Guide for your product.

Before you install SDDPCM, you must:

- Determine that you have the correct installation package
- Remove the SDD package, if it is installed.
- Remove the `ibm2105.rte` (version 32.6.100.x) and/or devices. `fc.disk.ibm.rte` (version 1.0.0.x), if they are installed.
- Install the AIX fibre-channel device drivers, if necessary.
- Verify and upgrade the fibre channel adapter firmware level
- Install the SDDPCM host attachment: `devices.fcp.disk.ibm.mpio.rte` (version 1.0.0.15 or later), or `devices.sas.disk.ibm.mpio.rte` (version 1.0.0.0 or later).

Determining the correct installation package

The package name of SDDPCM for different AIX operating system levels is the following:

- `devices.sddpcm.52.rte` for AIX 5.2
- `devices.sddpcm.53.rte` for AIX 5.3
- `devices.sddpcm.61.rte` for AIX 6.1

Determining if the SDD package is installed

To determine if the SDD is installed:

1. Use the `lslpp -l *ibmSdd*` and `lslpp -l devices.sdd*` commands to determine if any SDD package is installed on the system.
2. If SDD is installed for supported storage device configuration, you must unconfigure and remove all SDD vpath devices, and then uninstall the SDD package. See “Removing SDD from an AIX host system” on page 49.

Determining if the `ibm2105.rte` package is installed

To determine if the `ibm2105.rte` package is installed:

1. Use the `lslpp -l *ibm2105*` command to determine if any `ibm2105.rte` with VRMF 32.6.100.xx is installed.
2. If `ibm2105.rte` is installed, you must:
 - a. Unconfigure and remove all supported storage devices.

- b. Use **smitty** to uninstall the `ibm2105.rte` package. If `ibm2105.rte` is installed, you must remove all supported storage devices and uninstall the `ibm2105.rte` package. See “Removing SDD from an AIX host system” on page 49.

Determining if the `devices.fcp.disk.ibm.rte` package is installed

To determine if the `devices.fcp.disk.ibm.rte` package is installed:

Note: The following steps are not required for RSSM devices.

1. Use the `lslpp -l devices.fcp.disk.ibm*` command to determine if any `devices.fcp.disk.ibm.rte` with VRMF 1.0.0.x is installed.
2. If `devices.fcp.disk.ibm.rte` is installed, you must:
 - a. Unconfigure and remove all supported storage devices.
 - b. Use **smitty** to uninstall the `devices.fcp.disk.ibm.rte` package. If `devices.fcp.disk.ibm.rte` is installed, you must remove all supported storage devices and uninstall the `devices.fcp.disk.ibm.rte` package. See “Removing SDD from an AIX host system” on page 49.

Installing the AIX fibre-channel device drivers

Note: Installing AIX fibre-channel device drivers is not a requirement for RSSM devices.

You must check for the latest information on fibre-channel device driver APARs, maintenance-level fixes, and microcode updates at the following website:

www-933.ibm.com/support/fixcentral/

Complete the following steps to install the AIX fibre-channel device drivers from the AIX compact disk:

1. Log in as the root user.
2. Load the compact disc into the CD-ROM drive.
3. From your desktop window, enter **smitty install_update** and press **Enter** to go directly to the installation panels. The **Install and Update Software** menu is displayed.
4. Highlight **Install Software** and press **Enter**.
5. Press **F4** to display the **INPUT Device/Directory for Software** panel.
6. Select the compact disc drive that you are using for the installation; for example, `/dev/cd0`, and press **Enter**.
7. Press **Enter** again. The **Install Software** panel is displayed.
8. Highlight **Software to Install** and press **F4**. The **Software to Install** panel is displayed.
9. The fibre-channel device drivers include the following installation packages:

devices.pci.df1080f9

The adapter device driver for RS/6000 or IBM System p with feature code 6239.

devices.pci.df1000f9

The adapter device driver for RS/6000 or IBM System p with feature code 6228.

devices.pci.df1000f7

The adapter device driver for RS/6000 or IBM System p with feature code 6227.

devices.common.IBM.fc
The FCP protocol driver.

devices.fcp.disk
The FCP disk driver.

Select each one by highlighting it and pressing **F7**.

10. Press **Enter**. The **Install and Update from LATEST Available Software** panel is displayed with the name of the software you selected to install.
11. Check the default option settings to ensure that they are what you need.
12. Press **Enter** to install. SMIT responds with the following message:

```
+-----+
| ARE YOU SURE??
| Continuing may delete information you may want to keep. 413
| This is your last chance to stop before continuing. 415
+-----+
```

13. Press **Enter** to continue. The installation process can take several minutes to complete.
14. When the installation is complete, press **F10** to exit from SMIT. Remove the compact disc.
15. Check to see if the correct APARs are installed by entering the following command:

```
instfix -iv | grep IYnnnnn
```

where *nnnnn* represents the APAR numbers.

If the APARs are listed, that means that they are installed. If they are installed, go to “Configuring supported storage MPIIO-capable devices” on page 120. Otherwise, go to step 3.

16. Repeat steps 1 through 14 to install the APARs.

Verifying and upgrading the fibre-channel adapter firmware level

Use the following procedures to verify and upgrade your current fibre-channel adapter firmware level.

Note: The following procedures do not apply to RSSM devices.

Verifying the adapter firmware level:

You must verify that your current adapter firmware is at the latest level. If your current adapter firmware is not at the latest level, you must upgrade to a new adapter firmware (microcode). To check the current supported firmware level for fibre-channel adapters, go to the following website:

www-933.ibm.com/support/fixcentral/

Complete the following steps to verify the firmware level that is currently installed:

1. Enter the `lscfg -vl fcsN` command. The vital product data for the adapter is displayed.
2. Look at the **ZB** field. The **ZB** field should look similar to:

```
(ZB).....S2F3.30X1
```

To verify the firmware level, ignore the second character in the **ZB** field. In the example, the firmware level is sf330X1.

3. If the adapter firmware level is at the latest level, there is no need to upgrade; otherwise, the firmware level must be upgraded. To upgrade the firmware level, go to “Upgrading the fibre channel adapter firmware level.”

Upgrading the fibre channel adapter firmware level:

Upgrading the firmware level consists of downloading the firmware (microcode) from your AIX host system to the adapter. Before you upgrade the firmware, ensure that you have configured any fibre-channel-attached devices (see “Configuring fibre-channel-attached devices” on page 18). After the devices are configured, download the firmware from the AIX host system to the FCP adapter by performing the following steps:

1. Verify that the correct level of firmware is installed on your AIX host system. Go to the /etc/microcode directory and locate the file called df1000f7.XXXXXX for feature code 6227 and df1000f9.XXXXXX for feature code 6228, where XXXXXX is the level of the microcode. This file was copied into the /etc/microcode directory during the installation of the fibre-channel device drivers.
2. From the AIX command prompt, enter **diag** and press **Enter**.
3. Highlight the **Task Selection** option.
4. Highlight the **Download Microcode** option.
5. Press **Enter** to select all the fibre-channel adapters to which you want to download firmware. Press **F7**. The Download panel is displayed with one of the selected adapters highlighted. Press **Enter** to continue.
6. Highlight /etc/microcode and press **Enter**.
7. Follow the instructions that are displayed to download the firmware, one adapter at a time.

Fibre-channel HBA attributes for DS4000, DS5000, and DS3950 storage devices:

You must set the FC HBA **dyntrk** attribute to **yes** instead of the default setting, which is **no**.

Verifying and Upgrading SAS controller firmware

Refer to RSSM documentation, at the following URL, for instructions on how to verify and upgrade the integrated SAS controller firmware or SAS Expansion Card firmware.

www.ibm.com/systems/support/supportsite.wss/docdisplay?Indocid=MIGR-5078491&brandind=5000020

Installing and upgrading the AIX SDDPCM host attachment

The host attachment for SDDPCM adds supported storage devices information to allow AIX to properly configure these devices as MPIO-capable hdisks. The supported storage devices information allows AIX to perform the following actions:

- Identify the hdisks as 2105, 2145, 1750, 2107, 1820, DS4100, DS4200, DS4300, DS4500, DS4700, DS4800, DS5000, DS5020, or DS3950 hdisks.
- Set default hdisk attributes, such as queue_depth and timeout values.
- Indicate to the AIX device driver configure method to set 2105, 2145, 1750, 2107, 1820, DS3950, DS4100, DS4200, DS4300, DS4500, DS4700, DS4800, DS5000, or DS5020 hdisks as MPIO devices.

For fibre-channel attached devices, the AIX SDDPCM host attachment package name is `devices.fcp.disk.ibm.mpio.rte`. For RSSM devices, the AIX SDDPCM host attachment package name is `devices.sas.disk.ibm.mpio.rte`. See the SDDPCM Readme file for the latest host attachment level for the SDDPCM version that you are installing.

For AIX SDDPCM host attachment installation and upgrade instruction from the CD, see the *Host System Attachment Guide* for your storage device. Do not reboot the system before the SDDPCM package is installed on the system.

For AIX SDDPCM host attachment installation and upgrade from the SDD download Web site, use following procedure:

1. Download code from the website that is appropriate for your system.
2. Move the downloaded file to the `/usr/sys/inst.images` directory.
3. Change the directory to the `/usr/sys/inst.images` directory.
4. Untar the downloaded file. For example:

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

or

```
tar -xvf devices.sas.disk.ibm.mpio.rte.tar
```

5. Update the `.toc` file. For example:

```
pwd
rm -i .toc
inutoc .
grep -i devices.fcp.disk.ibm .toc
```

or

```
pwd
rm -i .toc
inutoc .
grep -i devices.sas.disk.ibm .toc
```

This command should reflect the newer SDDPCM host attachment version that will be uploaded.

6. From your desktop window, type **smitty install_update** and press **Enter** to go directly to the installation panels. The Install and Update Software menu is displayed.
7. Highlight **Install Software** and press **Enter**.
8. Type `.` to indicate the current directory and press **Enter**.
9. Highlight **Software to Install** and press **F4**. The Software to Install panel is displayed.
10. Select the `devices.fcp.disk.ibm.mpio.rte` package.
11. Press **Enter**. The Install and Update from the LATEST Available Software panel is displayed with the name of the software that you selected to install.
12. Check the default option settings to ensure that they are what you need.
13. Press **Enter** to install. SMIT responds with the following message:

```
+-----+
| ARE YOU SURE??                               |
| Continuing may delete information you may want to keep. |
| This is your last chance to stop before continuing.     |
+-----+
```

14. Press **Enter** to continue. The installation process may take a few minutes to complete.
15. When the installation or upgrade is complete, press **F10** to exit from SMIT.

- If this is a host attachment installation, do not reboot the system until you have the SDDPCM package installed. If this is a host attachment upgrade, and the SDDPCM package is already installed on the system, then reboot the system to complete the host attachment upgrade.

Attention: Do not reboot the system if you only have the SDDPCM host attachment package installed.

Installing and upgrading SDDPCM

Use these topics to install and upgrade SDDPCM.

Installing SDDPCM from CD-ROM

SDDPCM is released as an AIX installation image. The SDDPCM install image resides in the `/usr/sys/inst.images/SDDPCM` directory on CD-ROM directory. Because the package does not reside in the `/usr/sys/inst.images` directory, which is the default directory for the AIX install program, you must mount the CD-ROM file system before you can use SMIT to install SDDPCM from the CD-ROM directory.

Note:

- To mount the CD-ROM and install SDDPCM, you must have root access and AIX system administrator knowledge.
- The AIX SDDPCM host attachment package (`devices.fcp.disk.ibm.mpio.rte` or `devices.sas.disk.ibm.mpio.rte`) must be installed before you install the SDDPCM package (`devices.sddpcm.52.rte`, `devices.sddpcm.53.rte`, or `devices.sddpcm.61.rte`).

Creating and mounting the CD-ROM filesystem

To install SDDPCM from the CD-ROM, you must first create and mount the CD-ROM filesystem. Use SMIT to complete the following steps to create and mount the CD-ROM to CD-ROM file system.

Note: Throughout this procedure, `/dev/cd0` is used for the compact disc driver address. The driver address can be different in your environment.

- Log in as the root user
- Insert the compact disc into the CD-ROM drive.
- From your desktop window, enter **smitty fs** and press **Enter**.
- Select **Add / Change / Show / Delete File Systems** and press **Enter**.
- Select **CDROM File System** and press **Enter**.
- Select **Add a CDROM File System** and press **Enter**. The Add a CDROM File System panel is displayed.
- Select **DEVICE name** and select **F4**. The DEVICE name panel is displayed.
- Select the compact disc drive that you are using for the installation, (for example, `cd0`), and press **Enter**.
- Select **MOUNT POINT** and enter a directory where you want the CDROM File System to be mounted, (for example, `/cdmnt`).
- Click the default option settings for the other fields to ensure that they are what you need.

```
+-----+
+ Add a CDROM File System          +
+                                  +
+ Type or select values in entry fields.  +
+ Press Enter AFTER making all desired changes.  +
```

```

+
+
+ * DEVICE name          [Entry Fields] +
+ * MOUNT POINT          cd0             +
+ Mount AUTOMATICALLY at system restart?  [/cdmnt]      +
+ no                      +
+
+-----+

```

11. Press **Enter** to create the CDROM File System.
12. When the CDROM File System has been created, press **F10** to exit from smit.
13. From your desktop window, enter **smitty mount** and press **Enter**.
14. Select **Mount a File System** and press **Enter**. The Mount a File System panel is displayed.
15. Select **FILE SYSTEM name** and press **F4**
16. Select the CDROM File System that you created and press **Enter**.
17. Select **DIRECTORY** on which to mount and press **F4**.
18. Select the CDROM File System that you created and press **Enter**.
19. Select **TYPE of file system** and press **Enter**.
20. Select **cdrfs** as the type of file system and press **Enter**.
21. Select **Mount as a REMOVABLE file system?** and press **TAB** to change the entry to *yes*.
22. Select **Mount as a READ-ONLY system?** and press **TAB** to change entry to *yes*.
23. Click to check the default option settings for the other fields to ensure that they are what you need.

```

+-----+
+
+ Mount a File System
+ Type or select values in entry fields.
+ Press Enter AFTER making all desired changes.
+
+ [Entry Fields] +
+ FILE SYSTEM name  [/dev/cd0] +
+ DIRECTORY over which to mount  [/cdmnt] +
+ TYPE of file system  cdrfs +
+ FORCE the mount?     no +
+ REMOTE NODE containing the file system  [] +
+ to mount
+ Mount as a REMOVABLE file system?     yes +
+ Mount as a READ-ONLY system?         yes +
+ Disallow DEVICE access via this mount? no +
+ Disallow execution of SUID and sgid programs no +
+ in this file system?
+
+-----+

```

24. Press **Enter** to mount the file system.
25. When the file system has been mounted successfully, press **F10** to exit from smit.

Attention: Do not reboot the system if you only have the SDDPCM host attachment package installed.

Using the System Management Interface Tool facility to install SDDPCM

Use the System Management Interface Tool (SMIT) facility to install SDDPCM. The SMIT facility has two interfaces: nongraphical (enter **smitty** to invoke the nongraphical user interface) and graphical (enter **smit** to invoke the graphical user interface).

Throughout this SMIT procedure, `/dev/cd0` is used for the compact disc drive address. The drive address can be different in your environment. Complete the following SMIT steps to install the SDDPCM package on your system.

1. From your desktop window, `cd` to the directory where the CD-ROM file system is mounted, for example `/cdmnt`.
2. Go to the directory `usr/sys/inst.images/SDDPCM`.
3. From your desktop window, enter **smitty install_update** and press **Enter** to go directly to the installation panels. The Install and Update Software menu is displayed.
4. Highlight **Install Software** and press **Enter**.
5. Enter `.` to indicate the current directory and press **Enter**.
6. Highlight **Software to Install** and press **F4**. The Software to Install panel is displayed.
7. Select the `devices.sddpcm.52.rte`, `devices.sddpcm.53.rte`, or `devices.sddpcm.61.rte` package, based on the OS level.
8. Press **Enter**. The Install and Update from LATEST Available Software panel is displayed with the name of the software that you selected to install.
9. Check the default option settings to ensure that they are what you need.
10. Press **Enter** to install. SMIT responds with the following message:

```
ARE YOU SURE??  
Continuing may delete information you may want to keep.  
This is your last chance to stop before continuing.
```

11. Press **Enter** to continue. The installation process can take several minutes to complete.
12. When the installation is complete, press **F10** to exit from SMIT.

Unmounting the CD-ROM File System

After successfully installing SDDPCM, use the following procedure to unmount CD-ROM file system in order to remove the CD-ROM:

1. Go to the root (`/`) directory.
2. Enter **umount /cdmnt** and press **Enter** to unmount the CD-ROM file system from the `/cdmnt` directory.
3. Enter **rmfs /cdmnt** and press **Enter** to remove the CD-ROM file system.
4. Remove the CD-ROM.

Installing SDDPCM from downloaded code

Use this procedure to install SDDPCM from downloaded code:

1. Download code from the website that is appropriate for your system
2. Move the downloaded files to the `/usr/sys/inst.images` directory
3. Change directory to the `/usr/sys/inst.images` directory.
4. Untar the downloaded file: for example, `tar -xvf devices.sddpcm.53.rte.tar`
5. Update the `.toc` file. For example:

```
pwd  
rm -i .toc  
inutoc .  
grep -i sdd .toc
```

This command should reflect the newer SDDPCM code version that will be updated.

6. Continue the installation by following the instructions beginning in step 3 on page 113.

Installing SDDPCM with the AIX OS from an AIX NIM SPOT server to the client SAN boot disk or the internal boot disk

You can install SDDPCM from an AIX Network Installation Management (NIM) server to the client SAN boot disk or the internal boot disk at the same time that the AIX OS is installed. You must set up the NIM master and create the `lpp_source` and Shared Product Object Tree (SPOT) resources with the images on a file system, which is either NFS-exported or is obtained from a CD or DVD.

Prepare for the NIM SPOT installation with AIX OS and SDDPCM on the client's SAN boot disk or the internal boot disk. To do this, first set up a NIM master and create the `lpp_source` and SPOT resource. You can use the System Management Interface Tool (SMIT) facility to implement the following procedures:

1. Install the following filesets to set up the system as an NIM master:

```
bos.sysmgt.min.master
bos.sysmgt.nim.spot
```

2. Initialize the NIM master system by running the `smitty nim_config_env` command.
3. Create a new `lpp_source` and SPOT resource by running the `smitty nim_config_env` command.
4. Add the SDDPCM fileset to the newly created `lpp_source` by running the `smitty nim_task_inst` command.
5. Create a SPOT from the new `lpp_source` by running the `smitty nim_config_env` command.
6. Define an NIM client by running the `smitty nim` command.

See the NIM task roadmap on the Web for detailed information on how to complete these tasks:

publib16.boulder.ibm.com/pseries/en_US/aixins/insgdrf/nim_roadmap.htm#nim_roadmap

After you have successfully prepared for the NIM SPOT installation, you are ready to use the SMIT tool to start the NIM installation on the client system:

1. Run the `smitty nim` command.
 - a. Click **Perform NIM Administration Tasks > Manage Network Install Resource Allocation > Manage Machines > Allocate Network Install Resources**.
 - b. Select the hostname of the client that you defined previously.
 - c. Select the `lpp_source` and SPOT resources that you created previously, and then press **Enter**.
2. Run the `smitty nim` command again.
 - a. Click **Perform NIM Administration Tasks > Manage Machines > Perform Operations on Machines**.
 - b. Select the hostname of the client that you selected previously.
 - c. Click `bos_inst`.
 - d. Set the **ACCEPT new license agreements** field to **Yes**, and then press **Enter**.

The system automatically reboots after the **smitty nim** task completes. Use the following command to check the SAN boot disk and make sure the boot disk is configured with SDDPCM:

```
lsattr -El hdiskX (SAN boot disk device name)
```

From the output of this command, check the ODM attribute PCM to ensure that the value is *PCM/friend/sddpcm* or *PCM/friend/sddappcm*.

Updating SDDPCM

The following sections discuss methods of updating SDDPCM, verifying the currently installed version of SDDPCM, and the maximum number of devices that SDDPCM supports.

Updating SDDPCM packages by installing a newer base package or a program temporary fix

SDDPCM allows you to update SDDPCM by installing a newer base package or a program temporary fix (PTF). A PTF file has a file extension of .bff (for example, *devices.sddpcm.52.rte.2.1.0.1.bff*) and can either be applied or committed when it is installed. If the PTF is committed, the update to SDDPCM is permanent; to remove the PTF, you must uninstall SDDPCM. If the PTF is applied, you can choose to commit or to reject the PTF at a later time. If you decide to reject the PTF, you will not need to uninstall SDDPCM from the host system.

Note: If you are not upgrading the operating system, regardless of whether you have SAN boot devices, you can update SDDPCM packages by installing a newer base package or a program temporary fix. Otherwise, see “Migrating SDDPCM during an AIX OS upgrade with multipath SAN boot devices (on supported storage hdisks)” on page 119.

After applying the base package or the PTF, reboot the system. The SDDPCM server daemon should automatically start after restarting the system. If it does not start automatically, start the SDDPCM server daemon manually.

Use the SMIT facility to update SDDPCM. The SMIT facility has two interfaces, nongraphical (enter **smitty** to invoke the nongraphical user interface) and graphical (enter **smit** to invoke the GUI).

If the base package or PTF is on a CD-ROM, you must mount the CD file system. See “Creating and mounting the CD-ROM filesystem” on page 111 for directions on how to mount the CD file system. In the following procedure, */dev/cd0* is used for the CD drive address. The drive address can be different in your environment.

Complete the following SMIT steps to update the SDDPCM package on your system:

1. Log in as the root user.
2. Type **smitty install_update** and press **Enter** to go directly to the installation panels. The Install and Update Software menu is displayed.
3. Select **Install Software** and press **Enter**.
4. Press **F4** to display the INPUT Device/Directory for Software panel.
5. Select either a CD drive that you are using for the installation or a local directory where the packages reside; for example, */dev/cd0*, and press **Enter**.
6. Press **Enter** again. The Install Software panel is displayed.

7. Select **Software to Install** and press **F4**. The Software to Install panel is displayed.
8. Select the base package or the PTF package that you want to install.
9. Press **Enter**. The Install and Update from LATEST Available Software panel is displayed with the name of the software that you selected to install.
10. If you only want to apply the PTF, select **Commit software Updates?** and tab to change the entry to *no*. The default setting is to commit the PTF. If you specify *no* to Commit Software Updates?, ensure that you specify *yes* to Save Replaced Files?.
11. Check the other default option settings to ensure that they are what you need.
12. Press **Enter** to install. SMIT responds with the following message:

```

+-----+
|ARE YOU SURE??
|Continuing may delete information you may want to keep.
|This is your last chance to stop before continuing.
+-----+

```

13. Press **Enter** to continue. The installation process can take several minutes to complete.
14. When the installation is complete, press **F10** to exit from SMIT.
15. Unmount the CD-ROM file system and remove the compact disc.

Committing or rejecting a program temporary fix update

Before you reject a PTF update, you need to unconfigure and remove all supported storage devices from your host system. Committing a PTF does not require this extra step. Complete the following steps to commit or reject a PTF update with the SMIT facility. The SMIT facility has two interfaces: nongraphical (enter **smitty** to invoke the nongraphical user interface) and graphical (enter **smit** to invoke the GUI).

1. Log in as the root user.
2. From your desktop window, enter **smitty install** and press **Enter** to go directly to the installation panels. The Software Installation and Maintenance menu is displayed.
3. Select **Software Maintenance and Utilities** and press **Enter**.
4. Select **Commit Applied Software Updates** to commit the PTF or select **Reject Applied Software Updates** to reject the PTF.
5. Press **Enter**. The Commit Applied Software Updates panel is displayed or the Reject Applied Software Updates panel is displayed.
6. Select **Software name** and press **F4**. The software name panel is displayed.
7. Select the Software package that you want to commit or reject.
8. Check the default option settings to ensure that they are what you need.
9. Press **Enter**. SMIT responds with the following message:

```

+-----+
|ARE YOU SURE??
|Continuing may delete information you may want to keep.
|This is your last chance to stop before continuing.
+-----+

```

10. Press **Enter** to continue. The commit or reject process can take several minutes to complete.
11. When the installation is complete, press **F10** to exit from SMIT.

Note: You do not need to restart the system even though the bosboot message may indicate that a restart is necessary.

Verifying the currently installed version of SDDPCM

You can verify your currently-installed version of SDDPCM by issuing one of the following commands:

```
lslpp -l *sddpcm*
```

or

```
pcmpath query version
```

Note: Using the `pcmpath` command requires that at least one SDDPCM-supported MPIO device is configured on the system.

Maximum number of devices that SDDPCM supports

For SDDPCM 2.6.5.0 or later, a maximum of 65536 configured devices are supported and a maximum of 16 paths per device is supported. For earlier versions of SDDPCM, a maximum of 1200 configured devices are supported and a maximum of 16 paths per device is supported.

Thus, the maximum number of host adapter ports that are supported is 16. However, with the round robin or load balance path selection algorithms, configuring more than four paths per device might impact the I/O performance. Use the minimum number of paths necessary to achieve sufficient redundancy in the SAN environment. The recommended number of paths per device is four.

To support more than 600 storage device LUNs, system administrators must first determine whether the system has sufficient resources to support a large number of devices. For more information, see “Preparing your system to configure more than 600 supported storage devices or to handle a large amount of I/O after queue depth is disabled” on page 40.

For AIX blade servers in an IBM BladeCenter S Chassis connected to RSSM devices, refer to the RSSM documentation at the following URL for the maximum number of RSSM LUNs supported:

www.ibm.com/systems/support/supportsite.wss/docdisplay?Indocid=MIGR-5078491&brandind=5000020

For AIX 5.3, a single host should manage a maximum of 1024 devices when devices have been enabled for Open HyperSwap on the host, with 8 logical paths configured for each copy set in the session.

For AIX 6.1, a single host should manage a maximum of 1024 devices when devices have been enabled for Open HyperSwap on the host, with 16 logical paths configured for each copy set in the session.

Migrating SDDPCM

The following sections discuss the methods of migrating SDDPCM with and without SAN boot devices

- “Migrating the supported storage SAN boot device or nonboot volume group from AIX default PCM to SDDPCM” on page 118
- “Migrating from SDDPCM to the AIX default PCM or to SDD” on page 119
- “Migrating from SDD with SAN boot devices (on supported storage hdisks) to SDDPCM with multipath SAN boot devices” on page 119

Migrating the supported storage SAN boot device or nonboot volume group from AIX default PCM to SDDPCM

The default reserve policy of the AIX base PCM is a single-path policy, which is SCSI-2 reserve. The path selection algorithm is fail_over, which means only one path is opened at a time and that path made SCSI-2 reserve to the disk. All I/O is routed to this path. This reserve policy and path selection algorithm can cause problems if, after the SDDPCM packages are installed, you open the default PCM device before you restart the system. The default PCM device can be opened if:

- You build a volume group and file system with AIX default PCM devices, and leave the volume groups active and file system mounted
- You configure default PCM devices as the backing devices of a virtual target device when you are in a VIOS environment

After the system starts, you might see some paths in the INVALID state. The INVALID state means that the path failed to open. This is because the SCSI-2 reserve is not released during the system restart; thus, only the paths previously opened with SCSI-2 reserve can be opened and used for I/O after system restart. You might not see paths in the INVALID state if your system is at AIX 5.2 TL10 or later or at AIX 5.3 TL07 or later, or if you have IY83717 or IY83847 installed on your system. Instead, you might see a heavier I/O select count on one path. This is because the SCSI-2 reserve is not released during the system restart. Even if all the paths are allowed to be opened, only opened paths that previously made SCSI-2 reserve can actually be used for I/O.

If you have supported storage SAN boot devices that are configured with AIX default PCM, and the reserve policy is single_path (SCSI-2 reserve), switching the boot devices from AIX default PCM to SDDPCM might result in this reservation conflict situation. If you install an SDDPCM version earlier than 2.2.0.1, you must always run the **relbootrsv** command to release the SCSI-2 reserve on SAN boot devices after you install the SDDPCM host attachment package and the SDDPCM package. Run the **relbootrsv** command before you restart the system and then run the following command against the hdisks that are part of the rootvg to verify that they are no longer reserved.

```
# pcmquerypr -Vh /dev/hdisk6
connection type: fscsi0
open dev: /dev/hdisk6

Attempt to read reservation key...

Attempt to read registration keys...
Read Keys parameter
      Generation : 0
      Additional Length: 0
resrvpolicy= no_reserve
Reserve Key provided by current host = none (hex)02bbb003
Not reserved.
```

If you install SDDPCM version 2.2.0.1 or later, the SCSI-2 reserve on SAN boot devices is released automatically during system boot.

In a VIOS environment, reservation conflict problems can occur on a virtual target device that used to have a default PCM device as its backing device. To prevent this problem, perform one of the following actions:

- Switch from the AIX default PCM to SDDPCM before you use the AIX default PCM device as the backing device of a virtual target device.

- Before switching from the AIX default PCM to SDDPCM, put the virtual target devices in the **Define** state. This properly closes the AIX default PCM and releases the SCSI-2 reserve before migrating to SDDPCM.

Reservation conflict problems can also occur on nonboot volume groups. To prevent this problem, perform one of the following actions:

- Switch from the AIX default PCM to SDDPCM before you make any volume groups and file systems.
- To switch from the AIX default PCM to SDDPCM, you must unmount file systems and vary off the volume group of the AIX default PCM to release the SCSI-2 reserve on the volume group before system restart.
- Issue **relbootrsv VGname** to release the SCSI-2 reserve on the active, nonboot volume group devices before you restart the system.

Note: If you specify a *VGname* (volume group name), **relbootrsv** releases the SCSI-2 reserve of the specified non-SAN boot volume group; otherwise, it releases the SCSI-2 reserve of a SAN boot volume group (rootvg).

Migrating from SDDPCM to the AIX default PCM or to SDD

Note: If you have supported storage devices configured with SDDPCM as MPIO boot devices, you need to contact IBM Customer Support for migration from SDDPCM to the AIX default PCM. The following procedure is applicable only for a system without SDDPCM MPIO devices configured as the SAN boot device.

To migrate from SDDPCM to the AIX default PCM or to SDD, you must first unconfigure the devices, stop the SDDPCM server daemon, and then uninstall the SDDPCM package and the SDDPCM host attachment package. See “Removing SDDPCM from an AIX host system” on page 123 for directions on uninstalling SDDPCM. After you uninstall SDDPCM, you can then restart the system to migrate supported storage MPIO devices to the AIX default PCM. If you want to migrate supported storage devices to SDD devices, you must then install the supported storage device host attachment for SDD and the appropriate SDD package for your system. Then restart the system to configure the supported storage devices to SDD vpath devices.

Migrating from SDD with SAN boot devices (on supported storage hdisks) to SDDPCM with multipath SAN boot devices

If you have supported storage devices configured with SDD and there are SAN boot devices with supported storage hdisk devices, you need to contact IBM Customer Support for migration from SDD to SDDPCM.

Migrating SDDPCM during an AIX OS upgrade with multipath SAN boot devices (on supported storage hdisks)

SDDPCM provides different packages to match the AIX OS level. If an AIX system is going to be upgraded to a different OS level; for example, from AIX 5.3 to AIX 6.1, you need to install the corresponding SDDPCM package for that OS level.

If you want to upgrade AIX OS and there are SAN boot devices with SDDPCM supported storage hdisk devices, you need to contact IBM Customer Support for migration from SDDPCM during the OS upgrade.

If you are not in SAN boot environment, or you are only upgrading the AIX OS Technology Level or Server Pack; for example, from AIX 5.3 TL04 to AIX 5.3 TL06, you can follow the procedures in “Updating SDDPCM” on page 115.

Configuring and unconfiguring supported storage MPIO-capable devices

After installing the MPIO-supported storage device host attachment and the SDDPCM package, you need to restart the system in order to configure supported storage devices as MPIO-capable devices. After the first system restart, you can then use the normal AIX command line configure programs to configure and unconfigure supported storage MPIO-capable devices.

After the system restarts, the SDDPCM server daemon (`pcmsrv`) should automatically start.

Configuring supported storage MPIO-capable devices

The newly installed supported storage devices must be configured as MPIO-capable devices before you can use them. Use one of the following commands to configure these devices:

- `cfgmgr` command

Note: If operating in a switched environment, issue the `cfgmgr` command once for each host adapter each time a device is added.

If you use the `cfgmgr` command to configure supported storage MPIO devices, you might need to start the SDDPCM server daemon manually, if it has not already started. See “SDDPCM server daemon” on page 139 for information describing how to check the daemon status and how to manually start the daemon.

- `shutdown -rF` command to restart the system.

After the system restarts, the SDDPCM server daemon (`pcmsrv`) should automatically start.

Verifying the SDDPCM configuration

To verify the SDDPCM configuration, you can use one of the following:

- SMIT MPIO management submenu
- SDDPCM `pcmpath query device` command

Complete the following steps using SMIT to verify the SDDPCM configuration on an AIX host system:

1. Enter `smitty MPIO` from your desktop window. The **MPIO management** menu is displayed.
2. Select MPIO Device Management and press **Enter**. The MPIO Device Management panel is displayed.
3. Select List ALL MPIO Devices and press **Enter**. All MPIO devices on the host are listed.
4. Search for all IBM MPIO FC XXXX devices, where XXXX can be 2105, 2107, 1750, or SAN Volume Controller MPIO devices. Also search for MPIO FC 2145 if you have a host attachment package for SDDPCM with version 1.0.0.8 or later. Ensure that these devices are in the **Available** state.

Search for all IBM MPIO DSXXXX array disks, where XXXX can be 4100, 4200, 4300, 4400, 4700, 4800, 5000, 5020, or 3950 devices. Ensure that these devices are in the **Available** state

Search for all IBM MPIO SAS 1820 devices. Ensure that these devices are in the **Available** state.

You can also use the SDDPCM **pcmpath query device** command to query the configuration status of supported storage devices.

Note:

1. If none of the supported storage devices are configured successfully as MPIO devices, the **pcmpath query device** command will fail.
2. Because the AIX disk-configure method supports parallel device configuration, the same MPIO logical device name can be configured with different supported storage devices after each device configuration. If you have a program or script that has the MPIO logical device name hardcoded, ensure that you check the MPIO logical device name binding with the storage device LUN. If there is a change to the binding between an MPIO logical device name and the storage device LUN, you must revise your program or script.

Dynamically adding and removing paths or adapters

SDDPCM supports dynamically adding or removing extra paths or adapters while I/O is running. To add extra paths that are attached to a fibre-channel adapter or a SAS adapter, to existing available devices, enter:

```
| cfgmgr -vl fscsiX
```

or

```
| cfgmgr -vl sasX
```

When the command returns successfully, the paths are added to the devices.

To check the device configuration status, enter:

```
lspath -l hdiskX -H -F "name path_id parent connection status"
```

or

```
pcmpath query device X
```

To add a new fibre-channel adapter to existing, available, supported, storage MPIO devices, enter:

```
cfgmgr -vl fscX
```

To add a new SAS adapter, depending on the parent driver, enter:

```
cfgmgr -vl mptsasX
```

or

```
cfgmgr -vl sissasX
```

To check the adapter configuration status, enter:

```
pcmpath query adapter
```

or

pcmpath query device

To dynamically remove all paths under a parent fibre-channel adapter from a supported storage MPIO device, enter:

```
rmpath -dl hdiskX -p fscsiY
```

To dynamically remove a fibre-channel or SAS adapter and all child devices from supported storage MPIO devices, use **smit mpio**, or enter the following on the command line:

```
rmdev -l fscsiX -R
```

or

```
rmdev -l sasY -R
```

To dynamically remove a particular path, run **smit mpio**, or enter one of the following commands on the command line:

```
rmpath -l hdiskX -p fscsiY -w connection location code
```

or

```
rmpath -l hdiskX -p sasY -w connection location code
```

or

```
rmpath -dl hdiskX -p fscsiY -w connection location code
```

or

```
rmpath -dl hdiskX -p sasY -w connection location code
```

Issue the following command to get a particular path connection location code:

```
lspath -l hdiskX -H -F "name path_id parent connection status"
```

Note: You cannot remove the last path from a supported storage MPIO device. The command will fail if you try to remove the last path from a supported storage MPIO device.

Unconfiguring supported storage MPIO-capable devices

Note: Contact IBM support if your system has SAN boot MPIO devices configured with SDDPCM on supported storage system disks. The SAN boot MPIO devices cannot be unconfigured with following procedures.

To remove all non-SAN supported storage MPIO-capable devices:

1. Unmount the file systems of all supported storage devices.
2. Vary off all supported storage device volume groups.
3. Stop **pcmsrv**. See "Stopping the SDDPCM server" on page 140 for the procedure.

4. For fibre-channel-attached devices, enter the following command for each adapter:

```
rmdev -dl fcsX -R
```

5. For RSSM devices, enter the following command:

```
lsdev -Cc disk | grep 1820 | awk '{print $1}' | xargs -n1 rmdev -dl
```

Switching DS4000, DS5000, or DS3950 storage device configurations

For SDDPCM version 2.5.1.0 and subsequent releases, with systems having AIX 5.3 TL08 and AIX 6/1 TL02 and required AIX APAPs installed, users can use the **AIX manage_disk_drivers** command to display the configured drivers for DS3950, DS4100, DS4200, DS4300, DS4500, DS4700, DS4800, DS5100, and DS5020 storage devices, and a list of supported drivers for these device models. Supported drivers are SDDPCM (displayed as AIX_SDDAPPCM), AIX native PCM (displayed as AIX_APPPCM), and AIX_FCP array driver (displayed as AIX_fcpcarray; only supports DS4K storage devices). For example:

```
[root@telephone] /home/root # manage_disk_drivers -l
```

```
Device Present Driver Driver Options
2810XIV AIX_AAPCM AIX_AAPCM,AIX_non_MPIO
DS4100 AIX_SDDAPPCM AIX_APPPCM,AIX_fcpcarray
DS4200 AIX_SDDAPPCM AIX_APPPCM,AIX_fcpcarray
DS4300 AIX_SDDAPPCM AIX_APPPCM,AIX_fcpcarray
DS4500 AIX_SDDAPPCM AIX_APPPCM,AIX_fcpcarray
DS4700 AIX_SDDAPPCM AIX_APPPCM,AIX_fcpcarray
DS4800 AIX_SDDAPPCM AIX_APPPCM,AIX_fcpcarray
DS3950 AIX_SDDAPPCM AIX_APPPCM
DS5020 AIX_SDDAPPCM AIX_APPPCM
DS5100 AIX_SDDAPPCM AIX_APPPCM
```

This AIX command also allows users to switch DS4000, DS5000, and DS3950 storage device configuration drivers among these driver options. However, switching device configuration drivers requires a system reboot to reconfigure the devices after running this command. The syntax of this command function is:

```
manage_disk_drivers -d device -o driver_option
```

Removing SDDPCM from an AIX host system

Before you remove the SDDPCM package from your AIX host system, all supported storage devices must be unconfigured and removed from your host system, or migrated to the AIX default PCM. The SDDPCM server daemon must be stopped.

Note:

1. With version 3.0.0.0, it's no longer required to manually stop the SDDPCM server daemon (**pcmsrv**).
2. SDDPCM supports MPIO-supported storage devices as SAN boot devices. If your server has a supported storage device configured with SDDPCM as the SAN boot device, contact IBM support for your SAN boot configuration environment.

To remove SDDPCM, you must remove both the host attachment for SDDPCM and the SDDPCM software packages before you reconfigure supported storage devices

or restart the system. Otherwise, the devices can be in the *Defined* state and cannot be configured as either MPIO or non-MPIO devices.

After all the supported storage devices are removed or migrated to the AIX default PCM and the SDDPCM server daemon (**pcmsrv**) is stopped, complete the following steps to remove the SDDPCM software package:

1. Enter **smitty deinstall** from your desktop window to go directly to the Remove Installed Software panel.
2. Press **F4** in the SOFTWARE name field to bring up a list of packages and press the **F7** key to select the package to uninstall.
3. Press **Tab** in the **PREVIEW Only?** field to toggle between Yes and No. Select **No** to remove the software package from your AIX host system.

Note: If you select **Yes**, the process stops at this point and previews what you are removing. The results of your precheck are displayed without removing the software. If the state for any supported storage MPIO device is either *Available* or *Defined*, the process fails.

4. Select **No** for the remaining fields on this panel.
5. Press **Enter**. SMIT responds with the following message:

```
ARE YOU SURE??
Continuing may delete information you may want to keep.
This is your last chance to stop before continuing.
```

6. Press **Enter** to begin the removal process. This might take a few minutes.
7. When the process is complete, the SDDPCM software package and the supported storage device host attachment for SDDPCM are removed from your system.

SDDPCM support for HACMP with Enhanced Concurrent Mode volume groups

Beginning with SDDPCM 2.1.2.0, SDDPCM supports HACMP V5.1, V5.3, and V5.4 on an AIX 5.2 TL07 (or later) and AIX 5.3 TL03 (or later) system with both concurrent and nonconcurrent resource groups. SDDPCM 2.4.0.0 supports HACMP V5.3, and V5.4 on an AIX 6.1 system with both concurrent and nonconcurrent resource groups.

This support requires definition of the shared volume groups as Enhanced Concurrent Mode volume groups, in either concurrent or nonconcurrent resource groups. This means that no reserve needs to be broken in response to a node failure, and hence any requirement on breaking reserves is removed. A special interaction between HACMP and LVM ensures that if the volume group is used in a nonconcurrent resource group, applications are allowed to access it on one node at a time. Only `no_reserve` policy is supported in both concurrent and nonconcurrent resource groups.

The Enhanced Concurrent Mode volume groups are sufficient to ensure high availability. However, if system ECM volume groups are in nonconcurrent resource groups, you should configure your SAN using the following guidelines:

- The interaction between HACMP and LVM to ensure that only one node has nonconcurrent access at a time is *advisory locking*. This is in contrast to the

mandatory locking provided by SCSI reserves. To ensure that production data is not inadvertently modified by nodes that are not in the HACMP cluster, the following should be done:

1. Use either physical cabling or zoning to ensure that only HACMP nodes have access to the shared LUNs. That is, non-HACMP nodes should be prevented by hardware from accessing the shared LUNs.
 2. Start HACMP on the cluster nodes at boot time. This ensures that HACMP will activate the appropriate access controls on the shared disk before applications have a chance to modify the access controls.
- Configure disk heartbeating to reduce the likelihood of one node considering the other *dead* and attempting to take over the shared disks. (This is known as a *partitioned cluster*, or *split brain syndrome*). If the shared disks consist of multiple enclosures, use one disk in each enclosure as a heartbeat path.

Different storage systems or models might support different versions of HACMP and SDDPCM. For information, see the interoperability matrix for your storage:

www.ibm.com/systems/storage/disk/ds6000/
www.ibm.com/systems/storage/disk/ds8000/
www.ibm.com/systems/storage/software/virtualization/svc/

For HACMP v5.1, v5.2, v5.3, and v5.4 for AIX5L support information, go to the following website:

publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.hacmp.doc/hacmpbooks.html

For HACMP up-to-date APAR information, go to the following website:

www14.software.ibm.com/webapp/set2/sas/f/hacmp/download/aix53.html

MPIO reserve policies

SDDPCM supports four types of MPIO reserve policies. You can select one of the four reserve policies based on their configuration environment or application needs. The supported reserve policies are:

- No Reserve reservation policy
- Exclusive Host Access Single Path reservation policy
- Persistent Reserve Exclusive Host Access reservation policy
- Persistent Reserve Shared Host Access reservation policy

No Reserve reservation policy

If you set MPIO devices with this reserve policy, there is no reserve being made on MPIO devices. A device without reservation can be accessed by any initiators at any time. I/O can be sent from all the paths of the MPIO device. This is the default reserve policy of SDDPCM.

Exclusive Host Access single-path reservation policy

This is the SCSI-2 reservation policy. If you set this reserve policy for MPIO devices, only the fail_over path selection algorithm can be selected for the devices. With this reservation policy, an MPIO device has all paths being opened; however, only one path made a SCSI-2 reservation on the device. I/O can only be sent

through this path. When this path is broken, reserve is released, another path is selected, and SCSI-2 reserve is reissued by this new path. All input and output is now routed to this new path.

Persistent Reserve Exclusive Host Access reservation policy

If you set an MPIO device with this persistent reserve policy, a persistent reservation is made on this device with a persistent reserve (PR) key. Any initiators who register with the same PR key can access this device. Normally, you should pick a unique PR key for a server. Different servers should have different unique PR key. I/O is routed to all paths of the MPIO device, because all paths of an MPIO device are registered with the same PR key. In a nonconcurrent clustering environment, such as HACMP, this is the reserve policy that you should select.

Current[®] HACMP clustering software supports no_reserve policy with Enhanced Concurrent Mode volume group. HACMP support for persistent reserve policies for supported storage MPIO devices is not available.

Persistent Reserve Shared Host Access reservation policy

If you set an MPIO device with this persistent reserve policy, a persistent reservation is made on this device with a persistent reserve (PR) key. However, any initiators that implemented persistent registration can access this MPIO device, even if the initiators are registered with different PR keys. In a concurrent clustering environment, such as HACMP, this is the reserve policy that you should select for sharing resources among multiple servers.

Current HACMP clustering software supports no_reserve policy with Enhanced Concurrent Mode volume group. HACMP support for persistent reserve policies for supported storage MPIO devices is not available.

SDDPCM active/passive storage device controller health-check feature

This controller health-check function, when enabled, is invoked when an enabled path become unavailable due to transport problems. Enabling this feature results in faster I/O failure time when the access to a LUN is lost. The faster failure time is controlled by the controller health check attribute setting: **cntl_delay_time**.

By default, the controller health-check feature is disabled. To enable this feature, set the following two ODM attributes to a non-zero value for the active/passive storage devices: **cntl_delay_time** and **cntl_hcheck_int**.

Note: Setting either value to zero (0) disables this feature. Setting **cntl_delay_time=1** also disables this feature.

For example, if you want to allow the storage device to come back on the fabric within 30 seconds after leaving the fabric, set **cntl_delay_time=30** and **cntl_hcheck_int=2**. The device, /dev/hdisk#, must not be in use when setting the ODM values with the **chdev** command, or the **chdev -P** option must be used, which requires a reboot to make the change to take effect.

Two **pcmpath** commands allow you to dynamically change these two controller health-check ODM attributes. See “**pcmpath set device cntlhc_delay**” on page 170 and “**pcmpath set device cntlhc_interval**” on page 170 for information about these commands.

CAUTION:

There are cases where the storage device may reboot both of the controllers and become inaccessible for a period of time. If the controller health check feature is enabled, and the `cntl_delay_time` is set too short, it may result in an I/O failure. Ensure that you have a mirrored volume with failover protection, or with a GPFS configuration, if you are running with controller health check enabled and the `cntl_delay_time` setting is under 60 seconds, or is not long enough to cover the temporary device inaccessible conditions that can occur during the storage concurrent code load, or other error injection operations.

SDDPCM ODM attribute settings

The following sections discuss the SDDPCM ODM attribute default settings, and how to change the attributes of the supported storage MPIO devices:

- “SDDPCM ODM attribute default settings”
- “Changing device reserve policies”
- “Changing the path selection algorithm” on page 128
- “Changing SDDPCM path healthcheck mode” on page 130
- “Changing SDDPCM path healthcheck time interval” on page 131
- “Changing the Open HyperSwap quiesce expire time” on page 131
- “Changing SDDPCM controller healthcheck delay_time” on page 133
- “Changing SDDPCM controller healthcheck interval time” on page 134

Note: For RSSM devices, refer to RSSM documentation, at the following URL, for information about setting `queue_depth` attribute

www.ibm.com/systems/support/supportsite.wss/docdisplay?lnodocid=MIGR-5078491&brandind=5000020

SDDPCM ODM attribute default settings

SDDPCM has following default attributes settings:

Attribute	Default value
device reserve policy	<code>no_reserve</code>
path selection algorithm	<code>load balance</code>
healthcheck mode	<code>nonactive</code>
healthcheck time interval	<code>60 seconds</code>
quiesce expire time	<code>120 seconds</code>

Changing device reserve policies

Use the `chdev` command to change the reserve policy for a device. Because `chdev` requires the device to be unconfigured and reconfigured, this is a disruptive operation.

The following reserve policies can be used with any of the supported path selection algorithms (see “Supported SDDPCM features” on page 101):

- no_reserve
- persistent reserve exclusive host access
- persistent reserve shared host access

When the reserve policy of a device is exclusive host access single path (SCSI-2), the only supported path selection algorithm is the fail_over algorithm. The fail_over algorithm selects one path at a time for all I/Os. When the active path fails, then an alternative path is selected. The SCSI-2 reserve is reissued by this alternative path

To change the device reserve policy to no_reserve, enter:

```
chdev -l hdiskX -a reserve_policy=no_reserve
```

If you want to change the reserve policy to one of the persistent reserve policies, you must provide a persistent reserve key at the same time that you change the device policy to one of the persistent reserve types. For example, to change the reserve policy to PR_shared:

```
chdev -l hdiskX -a PR_key_value=0x1234 -a reserve_policy=PR_shared
```

Note: Beginning with SDDPCM 2.1.0.0, 2 persistent reserve tools are provided to manage supported storage system MPIO devices. See “Persistent reserve command tools” on page 141 for more information.

Changing the path selection algorithm

Starting with SDDPCM 2.1.0.0, you can use the **pcmpath set device algorithm** command to dynamically change the path selection algorithm. See “pcmpath set device algorithm” on page 167 for information about this command.

You can also use the **chdev** command to change the path selection algorithm of a device. Because **chdev** requires that the device be unconfigured and then reconfigured, this is a disruptive operation.

Use the following command to change the device path selection algorithm to round robin:

```
chdev -l hdiskX -a algorithm=round_robin
```

You can change the reserve_policy and algorithm for a device with one command. For example, to change the reserve policy to no_reserve and the path selection algorithm to round robin:

```
chdev -l hdiskX -a reserve_policy=no_reserve -a algorithm=round_robin
```

Using the load balancing port algorithm

Starting with SDDPCM 2.4.0.0, you can use the load balancing port algorithm to achieve better balanced I/O performance in certain specific SAN configuration environments. This path selection algorithm supports only 2105, 2107, 1750, and 2145 storage devices. This algorithm is based on the existing load balancing algorithm, but it also uses the incorporated I/O statistics on each target port to

which the host is connected. You can use this algorithm for SAN configurations where I/O throughput is unbalanced on the storage targets.

The following examples demonstrate the possible configurations in which this algorithm can help to improve performance.

In Figure 4, link 2 between the switch and the storage is receiving twice the load of the other links. This condition creates a throughput imbalance on link 2 that prevents optimal load balancing on both hosts.

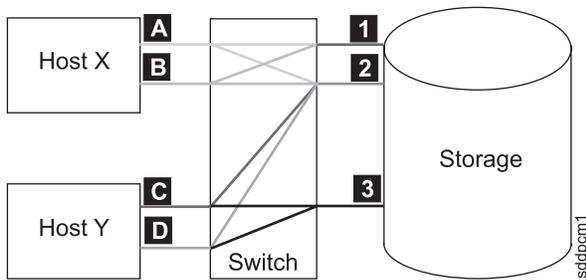


Figure 4. Workload imbalance when one link receives twice the load of the other links

Host X is attempting to balance the load across four paths (A1, A2, B1, B2). Host Y is attempting to balance the load across four paths (C2, C3, D2, D3).

In Figure 5, link 2 between the switch and the storage is more heavily loaded than link 1. This condition creates a throughput imbalance on link 2 that prevents optimal load balancing on Host X.

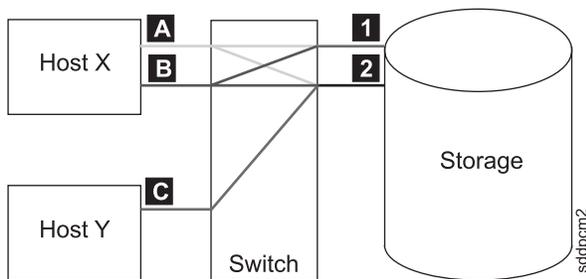


Figure 5. Workload imbalance when one link is more heavily loaded than another link

Host X is attempting to balance the load across four paths (A1, A2, B1, B2). Host Y has only one active path to the storage (C2).

In Figure 6 on page 130, one path is lost from a failed component, which results in the same workload imbalance shown in Figure 5.

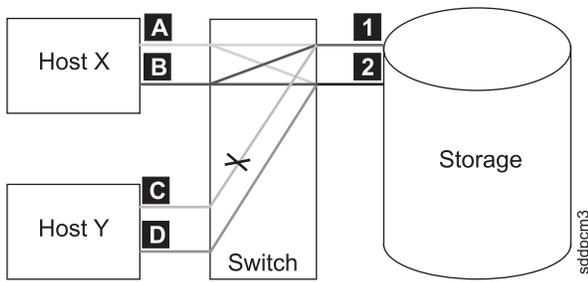


Figure 6. Workload imbalance when one host sharing workload across two paths loses one path

Host X is attempting to balance the load across four paths (A1, A2, B1, B2). Host Y has two paths to the storage, but only one path is active (C1, D2).

You can either use `pcmpath set device algorithm lbp` or `chdev -l hdiskX -a algorithm=load_balance_port` to configure the device with the load balancing port algorithm. To ensure that the load balancing port algorithm produces the desired result, all the hdisks on a host system that is attached to a given storage device must be configured with the same algorithm.

Note: This algorithm improves performance by balancing drastically unbalanced throughput on multiple target ports of the same storage. It cannot completely balance the throughput on the target ports because SDDPCM manages I/O load balancing only on the host where it is installed.

```
pcmpath query port target port number
pcmpath query portstats target port number
```

To display performance information about the target ports, you can use the `pcmpath query port target port number` command and the `pcmpath query portstats target port number` command. For more information, see “Using SDDPCM pcmpath commands” on page 145.

Changing SDDPCM path healthcheck mode

SDDPCM supports the path healthcheck function. If this function is enabled, SDDPCM tests opened paths and reclaims failed paths based on the value that is set in the following device healthcheck attribute:

```
hc_mode
```

Healthchecking supports the following modes of operations:

- Enabled - When this value is selected, the healthcheck command will be sent to paths that are opened with a normal path mode.
- Failed - When this value is selected, the healthcheck command is sent to paths that are in failed state.
- Nonactive - When this value is selected, the healthcheck command will be sent to paths that have no active I/O. This includes paths that are opened or in failed state.

If the algorithm selected is round robin or load balance, the healthcheck command will only be sent to failed paths, because the round robin and load-balancing algorithms route I/O to all opened paths that are functional. The default value setting of SDDPCM is *nonactive*.

Starting with SDDPCM 2.1.0.0, the **pcmpath set device hc_mode** command allows you to dynamically change the path healthcheck mode. See “pcmpath set device hc_mode” on page 169 for information about this command.

You can also use the **chdev** command to change the device path healthcheck mode. Because **chdev** requires that the device be unconfigured and then reconfigured, this is a disruptive operation. To change the path healthcheck mode to *failed*, issue the following command:

```
chdev -l hdiskX -a hc_mode=failed
```

Changing SDDPCM path healthcheck time interval

The `hc_interval` attribute will determine how often the paths of a device should be health-checked. The `hc_interval` attribute has a range of values from 0 - 3600 seconds. When a value of 0 is selected, the healthcheck function is disabled. Beginning with SDDPCM v2.1.2.3, the default value setting is changed from 20 (seconds) to 60 (seconds).

Starting with SDDPCM 2.1.0.0, the **pcmpath set device hc_interval** command allows you to dynamically change the path healthcheck time interval. See “pcmpath set device hc_interval” on page 168 for information about this command.

You can also use the **chdev** command to change the device path healthcheck time interval. Because **chdev** requires that the device be unconfigured and then reconfigured, this is a disruptive operation. To disable the path healthcheck interval function, issue the following command:

```
chdev -l hdiskX -a hc_interval=0
```

Note: Currently, the SDDPCM healthcheck function only checks the paths being opened. It does not healthcheck any path that is in the close state. The SDDPCM server daemon healthchecks `close_failed` paths. If a device healthcheck function is disabled, the SDDPCM server daemon healthchecks the device's failed path that is opened. See Chapter 9, “Using the SDDDSM or the SDDPCM daemon,” on page 365 for more information.

From SDDPCM v2.1.2.3 and later, a new feature of SDDPCM server daemon healthcheck function is introduced. The SDDPCM server daemon automatically starts or stops the healthcheck function on a device if you issue one of the following:

- **pcmpath set device m hc_interval 0** on the fly to disable a device internal healthcheck function
- **pcmpath set device m hc_interval n** to enable a device internal healthcheck function on the fly

Note: The SDDPCM server daemon only healthcheck FAILED paths. It does not implement healthcheck on opened or idle paths as the SDDPCM internal healthcheck does. This is the difference between the SDDPCM internal healthcheck and the SDDPCM server daemon healthcheck.

Changing the Open HyperSwap quiesce expire time

Beginning with SDDPCM 3.0.0.0, you can use the **chdev** command to change the Open HyperSwap quiesce expire time of a device. Because **chdev** requires that the

device be unconfigured and then reconfigured, this is a disruptive operation. Use the following command to change the device quiesce expire time:

```
chdev -l hdisk$i -a quiesce_timeout=n (30 < n < 600)
```

Note: Note: For a setup with a large number or the maximum number of devices, increase the **quiesce_timeout** to prevent a delayed recover command.

Supported AIX fibre-channel device driver features

SDDPCM supports the following fibre channel adapter driver features:

- Fast I/O failure of fibre channel devices
- Fibre-channel dynamic device tracking
- Changing SDDPCM controller healthcheck delay_time
- Changing SDDPCM controller healthcheck interval time

Fast I/O failure of fibre-channel devices

The AIX fibre channel adapter driver supports the new `fc_err_recov` attribute in AIX. If enabled, this attribute can reduce the amount of time needed to fail I/O in certain conditions. By default, this attribute has a value of `delayed_fail`, which means that the fast failure is disabled.

To enable fast failure, you must complete the following steps:

1. Change all the child devices of `fscsiX` on your system to the defined state by running the `rmdev -l fscsiX -R` command for all adapters on your system.
2. Enter the `chdev -l fscsiX -a fc_err_recov=fast_fail` command for all adapters on your system.
3. Enter the `cfgmgr` command to reconfigure all devices back to the available state.

You need to have the following Fibre Channel device driver APARs applied to your system in order to have all dynamic tracking and fast fail fixes:

APAR Description

IY37183

Dynamic Tracking and Fast I/O Failure of Fibre Channel Devices

IY44139

Dynamic Tracking: PERM I/O errors after Host HBA Cable Pull

IY44142

Dynamic Tracking: I/O Hang after Host HBA Cable Pull

IY44342

Dynamic Tracking: Ioctl call may fail after N_Port ID Change

IY44344

Dynamic Tracking: Back-to-Back Cable Move May Delay Error Recovery

IY45368

Fast Fail/Dynamic Tracking: FC Device Inaccessible after Move

IY46701

Dynamic Tracking & MPIO: Multiple Cable Swap Cause Path Failure

Note: You must set the FC HBA `fc_err_recov` attribute to `fast_fail` for DS4000, DS5000, and DS3950 storage devices.

Fibre channel dynamic device tracking

Dynamic Tracking support allows for the following scenarios without I/O errors or loss of access:

- Combine two switches in two different SANS into one SAN by connecting the switches with a cable and cascading the switches.
- Change the supported storage device FC cable from one switch port to a different switch port; Disconnected cable must be reconnected within 15 seconds.
- Swap two supported storage device FC cables on the switch ports. The disconnected cables must be reconnected within 15 seconds.
- Swap two supported storage device FC cables on the supported storage device ports. The disconnected cables must be reconnected within 15 seconds.

Note:

1. This 15 seconds window actually includes the time required to bring up the FC link after the cable is reconnected. For example, if it takes 4 seconds to bring up the FC link, the time allowed to leave the cable disconnected is 11 seconds. I/O will fail if the cable is disconnected longer than 11 seconds.
2. Fibre-channel device dynamic tracking does not support the case when a cable is moved from one adapter on the supported storage device to another free, previously unseen adapter on the supported storage device. The reason is that the world wide port name (WWPN) will be different for that adapter, and the supported storage device cannot track that adapter. The WWPN must be a static identifier of a remote port.

AIX has added a new attribute that must be changed to enable dynamic tracking. By default, dynamic tracking is disabled. To enable dynamic tracking, you must do the following:

1. Change all the child devices of fscsiX on your system to the defined state by issuing the **rmdev -l fscsiX -R** command for all adapters on your system.
2. Issue the **chdev -l fscsiX -a dyntrk=yes** command for all adapters on your system.
3. Issue **cfgmgr** to reconfigure all devices back to the available state.

Note: You must set the FC HBA **dyntrk** attribute to **yes** for DS4000, DS5000, and DS3950 storage devices

For more information about Dynamic Tracking and Fast I/O Failure, see:

www-01.ibm.com/support/docview.wss?rs=1203&context=SWG00&context=SWG10&context=SWGC0&context=HW182&dc=DB550&q1=dynamic+fast+tracking+%22and%22+Fast+I%2fO+failure&uid=sg1IY37183&loc=en_US&cs=UTF-8&lang=en

Installing APAR IY37183 also installs the file:

```
/usr/lpp/bos/README.FIBRE-CHANNEL
```

This file has more information about the Dynamic Tracking and Fast I/O Failure features.

Changing SDDPCM controller healthcheck delay_time

Starting with SDDPCM V2.5.2.0 and subsequent releases, the **cntl_delay_time** attribute is used to specify the amount of time, in seconds, that the active/passive

storage device's controller(s) will be health checked after a transport failure. At the end of the specified period, if no paths are detected as good, all pending and subsequent I/O to the device will fail until the device health checker detects that a failed path has recovered.

The **cntl_delay_time** attribute has a range of values from 0-300 seconds. When a value of 0 or 1 is selected, the controller healthcheck function is disabled. By default, this attribute value is set to 0, which means the controller health check function is disabled.

The **pcmpath set device cntlhc_delay** command allows you to dynamically change the controller health check delay_time. See “pcmpath set device cntlhc_delay” on page 170 for information about this command.

You can also use the **chdev** command to change the device controller healthcheck delay_time. Because chdev requires that the device be unconfigured and then reconfigured, this is a disruptive operation. Or you can use the **chdev -P** command, which requires a reboot to make the change to take effect.

Note: To enable the active/passive storage device controller health check function, the values of both controller health check attributes **cntl_delay_time** and **cntl_hcheck_int** must be set to a non-zero value.

Changing SDDPCM controller healthcheck interval time

Starting with SDDPCM V2.5.2.0 and subsequent releases, the **cntl_hcheck_int** attribute is used to specify the amount of time, in seconds, that the next controller health check command is issued. When the controller health check is enabled, the first controller health check is issued after a storage fabric transport failure is detected. This value must be less than the **cntl_delay_time** unless set to zero (0), to disable this feature.

The **cntl_delay_time** attribute has a range of values from 0-300 seconds. When a value of 0 is selected, the controller healthcheck function is disabled. By default, this attribute value is set to 0, which means the controller health check function is disabled.

The **pcmpath set device cntlhc_interval** command allows you to dynamically change the controller health check interval time. See “pcmpath set device cntlhc_interval” on page 170 for information about this command.

You can also use the **chdev** command to change the device controller healthcheck delay_time. Because chdev requires that the device be unconfigured and then reconfigured, this is a disruptive operation. Or you can use the **chdev -P** command, which requires a reboot to make the change to take effect.

For example, to enable the device controller health check function, issue the following command to set cntl_delay_time and cntl_hcheck_int to non-zero values:

```
chdev -l hdiskX -a cntl_delay_time=60 -a cntl_hcheck_int=3
```

Multipath SAN boot support

This section describes how to configure supported storage system MPIO devices as the SAN boot device.

Attention: The `sddpcm` default reserve policy is `no_reserve` and the default path selection algorithm is `load_balance`. If the SAN boot devices are configured with the default reserve policy and path selection algorithm, you must not configure these SAN boot devices on other AIX servers. This is because no `scsi-2` reserve or `scsi-3` persistent reserve is implemented on these SAN boot devices to prevent access from another AIX server.

Configuring supported storage system MPIO devices as the SAN boot device

A supported storage MPIO device can be used as the system boot device. To configure the supported storage device boot device with the SDDPCM module:

1. Select one or more supported storage system devices as the boot device.
2. Install one of the following AIX operating systems on the selected supported storage devices:
 - If the selected supported storage device is ESS, the required operating system is AIX 5.2 TL06 (or later), AIX 5.3 TL02 (or later), or AIX 6.1 TL0 (or later).
 - If the selected supported storage device is DS6000, the required operating system is AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1 TL0 (or later).
 - If the selected supported storage device is DS8000, the required operating system is AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1 TL0 (or later).
 - If the selected supported storage device is SAN Volume Controller, the required operating system is AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1 TL0 (or later).
 - If the selected supported storage device is DS4000, DS5000, or DS3950, the required operating system is AIX53 TL8 (or later) or AIX61 TL2 (or later).
 - If the selected supported storage device is RSSM, the required operating system is AIX 6.1 TL03 (or later). Refer to the following URL for RSSM documentation with instructions on how to install the AIX operating system on a JS blade with remote volumes:
www.ibm.com/systems/support/supportsite.wss/docdisplay?lnocid=MIGR-5078491&brandind=5000020
3. Restart the system. The supported storage boot device is configured as an MPIO device with AIX default PCM.

Note: For IBM DS4000 storage devices, if the OS is AIX53, or if the system was migrated from AIX53 to AIX61 (or later), the DS4000 devices are configured by the AIX FCPArray(RDAC) device driver instead of the AIX native MPIO PCM.

4. Install the supported storage device host attachment for SDDPCM and the SDDPCM packages.
5. To release `scsi-2` reserve on boot devices with SDDPCM v2.2.0.0 or earlier, you must issue the `relbootrsv` command. For boot devices with SDDPCM v2.2.0.1 or later, `scsi-2` reserve is automatically released during the system reboot. If you want to release non-rootvg `scsi-2` reserves, provide the volume group name as a parameter. For example: `relbootrsv vgname`.
6. Restart the system.

All supported storage MPIO devices, including supported storage MPIO SAN boot devices, are now configured with SDDPCM.

When you convert a boot device with SDDPCM v2.2.0.0 and earlier from the AIX default PCM to SDDPCM, you must issue the **relbootrsv** command, as shown in step 5. If you fail to do so, you might encounter a problem where either all paths of the boot device cannot be opened successfully or they can be opened but cannot be used for I/O. This problem occurs because the AIX default PCM has a default reserve policy of single-path (SCSI-2). See “Migrating the supported storage SAN boot device or nonboot volume group from AIX default PCM to SDDPCM” on page 118 for information about solving this problem.

There is a known problem during the SAN boot configuration. After the system is restarted followed by the operating system installation on the supported storage MPIO devices, you might see that some paths of the rootvg are in Failed path state. This can happen even if the system is successfully restarted.

This problem is corrected in AIX 5.2 TL08 or later and AIX 5.3 TL04 or later. Apply the following APARs on these OS levels after the first reboot followed by the operating system installation:

- AIX 5.2 TL08 or later: apply APAR IY83717
- AIX 5.3 TL04 or later: apply APAR IY83847

No APAR is available to correct this problem on AIX52 TL07 and AIX53 TL03. If you configure a SAN boot device with supported storage MPIO devices on one of these operating system levels and experience this problem, you can manually recover these failed paths by issuing one of the following commands:

- **chpath -s E -l hdiskX -p fscsiX**
- **pcmpath set device M path N online**

Support system dump device with the supported storage system MPIO device

You can choose a supported storage MPIO device to configure with the system primary and secondary dump devices. You can configure the system dump device with the supported SAN boot device, or with the nonboot device. The path selection algorithm for the system dump device automatically defaults to *failover_only* when the system dump starts.

During the system dump, only one path is selected for dump requests. If the first path fails, I/O is routed to the next path being selected.

You must apply all the APARs for AIX 5.2 TL08 and later, or AIX 5.3 TL04 and later.

Dynamically enabling and disabling paths or adapters

This section describes how to dynamically enable and disable paths and adapters.

Dynamically enabling or disabling a path

There are three ways to dynamically enable (place online) or disable (place offline) a path:

1. Use the following **pcmpath** commands to change the path state:
pcmpath set device M path N online

or

pcmpath set device M path N offline

2. Use the path control commands provided by AIX.

AIX 5.2 TL04 (or later) or AIX 5.3 TL01 (or later) provides several new path control commands. These commands can be used to configure or remove paths, change path state (enable or disable), and display the path current state.

Use the following AIX path command to change the path state of the fibre-channel or SAS-connected path:

```
chpath -l hdiskX -s E|D -p fscsiX -w "5005076300c99b0a,5200000000000000"
```

or

```
chpath -l hdisk1 -s E|D -p sasX -w 5005076b074012a0,1000000000000000
```

Note: To find the parent adapter and connect location of a path, use the following AIX path command:

```
lspath -l hdiskX -H -F "name path_id parent connection"
```

3. Use the smitty MPIO management submenu.
 - a. Enter **smitty MPIO** and press **Enter**. This displays the MPIO Management panel.
 - b. Select MPIO Path Management and press **Enter**. This displays the MPIO Path Management panel.
 - c. Select Enable paths or Disable paths to enable or disable paths.

Note:

1. If you issue the **pcmpath set device M path N offline** command to a DS4000 or a DS5000 device and the path to be set offline is the last active path, it causes the device to fail over to the alternate controller. This is a time-consuming action, and might take a few seconds to complete; thus, after the **pcmpath set device M path N offline** command returns, the device might not have completed the controller failover yet. You must issue the **pcmpath query device M** command to check the result and to ensure that this set offline command completed successfully. If the command completed successfully, you can see that the active paths and passive paths are swapped. This is an asynchronous event.
2. For Open HyperSwap devices, an attempt to set the last active path of the source volume offline by using the **pcmpath set device M path N offline** command will be rejected. When the device is closed, it's allowed to set the last active path of the source volume offline. However, when the device is opened with all paths to the source volume offline, an unplanned HyperSwap will be triggered.

Dynamically enabling or disabling an adapter

The SDDPCM **pcmpath** command can be used to enable (place online) or disable (place offline) an adapter.

To disable an adapter, use the following command:

```
pcmpath set adapter N offline
```

Note:

1. Beginning with SDDPCM version 2.4.0.0, storage devices DS4000 and DS5000 are supported. These are active/passive storage devices, and behave differently

than ESS, DS6000, DS8000, RSSM, and SAN Volume Controller storage devices. To support the new devices, the syntax of the **pcmpath set adapter N online | offline** command changed by adding the extra parameter **aa | ap**. To set the adapter that connects to ESS, DS6000, DS8000, RSSM, or SAN Volume Controller storage devices to online or offline, enter the **aa** parameter. To set the adapter that connects to DS4000, DS5000, or DS3950 storage devices to online or offline, enter the **ap** parameter.

2. SDDPCM reserves the last path of a device. This command will fail if there is any device using the last path attached to this adapter.
3. For DS4000 and DS5000 storage devices, this set adapter offline command might cause multiple devices to fail over to an alternate controller if the device path that is attached to this adapter is the last active path. Furthermore, if some devices are owned by controller A and some devices are owned by controller B, the set adapter offline command might cause some devices to fail over from controller A to controller B and some devices to fail over from controller B to controller A. It is possible that the **pcmpath set adapter M offline** command returns a failure message, but some number of devices are failed over successfully and some devices are not. You must issue the **pcmpath query device** command to find out if the command is fully completed or partially completed.
4. For Open HyperSwap devices, when the device is open, an attempt to set the last active path of the source volume offline by the **pcmpath set adapter** command will be rejected. When the device is closed, it's allowed to set the last path of the source volume offline. However, when the device is opened with all paths to the source volume offline, an unplanned HyperSwap will be triggered.

Attention: If the AIX server has more than one DS4000, DS5000, or DS3950 device model configured, the **pcmpath set adapter N offline** command might cause multiple device models to fail over to the alternate controllers. If this is the case, the **pcmpath set adapter N offline** command is not supported.

Using the SDDPCM trace function

SDDPCM supports AIX trace function. The SDDPCM trace ID is 5A7. Trace ID 5A7 traces routine entry, exit, and error paths of the algorithm. To use it, manually turn on the trace function before the program starts to run, then turn off the trace function either after the program stops, or any time you need to read the trace report.

To start the trace function, enter:

```
trace -a -j 5A7
```

To stop the trace function, enter:

```
trcstop
```

To read the report, enter:

```
trcrpt | pg
```

To save the trace data to a file, enter:

```
trcrpt > filename
```

Note: To perform the SDDPCM trace function, you must have the `bos.sysmgt.trace` installation package installed on your system.

SDDPCM server daemon

The SDDPCM server (also referred to as *pcmsrv*) is an integrated component of SDDPCM 2.1.0.0 (or later). This component consists of a UNIX application daemon that is installed in addition to the SDDPCM path control module. The SDDPCM server daemon provides path-recovery function for SDDPCM devices and the First Time Data Capture function. For detailed information about the SDDPCM server daemon, see Chapter 9, “Using the SDDDSM or the SDDPCM daemon,” on page 365.

Verifying if the SDDPCM server has started

After you have installed SDDPCM and restarted the system, verify if the SDDPCM server (`pcmsrv`) has automatically started by entering `lssrc -s pcmsrv`. If `pcmsrv` has automatically started, the output from the `lssrc -s pcmsrv` command looks like this:

```
+-----+
| Subsystem  GROUP  PID  Status |
| pcmsrv          NNN  Active |
+-----+
```

where *NNN* is the process ID number.

The status of `pcmsrv` should be *Active* if the SDDPCM server has automatically started. If the SDDPCM server has not started, the status will be *Inoperative*. Go to “Starting the SDDPCM server manually” to proceed.

Because `pcmsrv` is bound to the SDDPCM kernel extension module, `pcmsrv` can fail to start if the SDDPCM is installed and the supported storage MPIO devices have not been configured yet. In this case, you can either restart the system or you can start `pcmsrv` manually after supported storage MPIO devices are configured.

Because `pcmsrv` is bound to the SDDPCM kernel extension module, in order to uninstall or upgrade SDDPCM, you must stop `pcmsrv` so that the SDDPCM kernel extension can be unloaded from the system. During an upgrade, the new SDDPCM kernel extension can be loaded into the system when supported storage MPIO devices are configured.

Starting the SDDPCM server manually

If `pcmsrv` did not start automatically after you performed the SDDPCM installation and configured the supported storage devices, you can start `pcmsrv` manually.

For SDDPCM 3.0.0.0 or later releases, you can start `pcmsrv` by entering:

```
startpcmsrv
```

For `sddpcm 2.6.0.1` or prior releases, you can start `pcmsrv` by entering:

```
startsrv -s pcmsrv -e XPG_SUS_ENV=ON
```

Go to “Verifying if the SDDPCM server has started” to see if you successfully started the SDDPCM server.

Stopping the SDDPCM server

You can manually stop **pcmsrv**.

For SDDPCM 3.0.0.0 or later releases, you can permanently stop **pcmsrv** by entering:

```
stoppcmsrv
```

For SDDPCM 2.6.0.1 or earlier releases, you can temporarily disable **pcmsrv** by entering:

```
stopsrc -s pcmsrv.
```

The **stopsrc -s pcmsrv** command will stop the current version of **pcmsrv**, but **pcmsrv** will start again if the system is restarted.

In certain cases, when you want to permanently disable **pcmsrv** on SDDPCM 2.6.0.1 or earlier releases, so that it does not start even after a system restart, you must remove the following line in the system init table (/etc/inittab):

```
srv:2:wait:/usr/bin/startsrc -s pcmsrv -e XPG_SUS_ENV=ON > /dev/null 2>&1
```

Attention: The SDDPCM server (**pcmsrv**) recovers failed paths for devices that have disabled the SDDPCM healthcheck function. Furthermore, SDDPCM server provides the First Time Data Capture function, which saves an SDDPCM internal trace into a log file. You should not stop **pcmsrv** permanently.

See Chapter 9, “Using the SDDDSM or the SDDPCM daemon,” on page 365 for information about the following tasks:

- Enabling or disabling a TCP/IP port for SDDPCM server (“Enabling or disabling the sddsrv or pcmsrv TCP/IP port” on page 369)
- Changing to a different port number for the SDDPCM server (“Changing the sddsrv or pcmsrv TCP/IP port number” on page 369)

AE daemon

For SDDPCM 3.0.0.0 or later releases, a UNIX[®] application daemon, AE server, is added to SDDPCM path control module. The AE server daemon interfaces with the TPC-R server and the SDDPCM kernel driver to provide OpenSwap functionality.

Verifying if the AE server has started

After you have installed SDDPCM and restarted the system, verify if the AE server has automatically started by entering **lssrc -s AE**. If AE has automatically started, the output from the **lssrc -s AE** command looks like this:

```
+-----+
Subsystem Group PID Status
AE NNN active
```

where NNN is the process ID number.

The status of AE should be *Active* if the AE server has automatically started. If the AE server has not started, the status will be *Inoperative*. Go to “Starting the AE server manually” on page 141 to proceed.

Starting the AE server manually

If AE server did not start automatically after you performed the SDDPCM installation, you can start AE by entering:

```
startAE
```

Go to “Verifying if the AE server has started” on page 140 to see if you successfully started the AE server.

Stopping the AE server manually

You can temporarily disable AE by entering:

```
stopAE
```

Attention: Don't manually stop AE while the application is running and there are devices configured. AE is an integral part of the Open HyperSwap functionality. If AE is not running, Open HyperSwap functionality is not available. It is important to ensure that AE is running.

SDDPCM utility programs

This section describes the persistent reserve command tools, the SDDPCM `pcmpath` commands, and a tool to query DS4000 and DS5000 storage subsystem and MPIO device configuration information.

Persistent reserve command tools

Starting with SDDPCM 2.1.0.0, SDDPCM supports two persistent reserve command tools. The following sections describe the persistent reserve command tools.

pcmquerypr

The `pcmquerypr` command provides a set of persistent reserve functions. This command supports the following persistent reserve service actions:

- Read persistent reservation key
- Release persistent reserve
- Preempt-abort persistent reserve
- Clear persistent reserve and registration keys

This command can be issued to all system MPIO devices, including MPIO devices not supported by SDDPCM.

The `pcmquerypr` command can be used in the following situation: the reserve policy of the SDDPCM MPIO devices is set to either persistent reserve exclusive host access (`PR_exclusive`) or persistent reserve shared host access (`PR_shared`), and persistent reserve has been left on the device by a node, blocking access by another node. The `pcmquerypr` command can be used in this situation to query, preempt, or clear the persistent reserve left by a node or server on the devices.

There are more cases when you might need this tool to solve persistent reserve related problems, such as unexpected persistent reserve left on devices because of failing to release the persistent reserve. Caution must be taken with the command, especially when implementing preempt-abort or clear persistent reserve service action. With preempt-abort service action not only the current persistent reserve key is preempted; it also aborts tasks on the LUN that originated from the

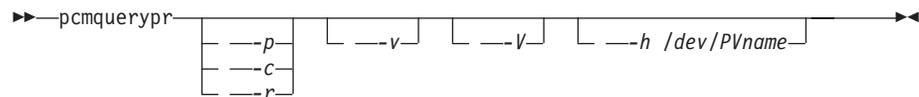
initiators that are registered with the preempted key. With clear service action, both persistent reservation and reservation key registrations are cleared from the device.

The following information describes in detail the syntax and examples of the **pcmquerypr** command.

Description

The **pcmquerypr** command implements certain SCSI-3 persistent reservation commands on a device. The device can be a supported storage MPIIO device. This command supports persistent reserve IN and OUT service actions, such as read, reservation key, release persistent reservation, preempt-abort persistent reservation, or clear persistent reservation and reservation key registrations.

Syntax



Flags:

- p** If the persistent reservation key on the device is different from the current host reservation key, the existing persistent reservation key on the device is preempted. This option can be issued only when the device is not already open.
- c** If there is a persistent reservation on the device, the persistent reservation is removed and all reservation key registrations on the device are cleared. This option can be issued only when the device is not already open.
- r** Removes the persistent reservation key on the device made by this host. This option can be issued only when the device is not already open.
- v** Displays the persistent reservation key if it exists on the device.
- V** Verbose mode. Prints detailed message.

Return code

If the command issued without options of **-p**, **-r** or **-c**, the command will return:

- 0** There is no persistent reservation key on the device, or the device is reserved by the current host
- 1** The persistent reservation key is different from the host reservation key
- 2** The command failed.

If the command issued with one of the options of **-p**, **-r** or **-c**, it returns:

- 0** The command was successful.
- 2** The command failed.

Examples

1. To query the persistent reservation on a device, enter **pcmquerypr -h /dev/hdisk30**.
This command queries the persistent reservation on the device without displaying. If there is a persistent reserve on a disk, it

returns 0 if the device is reserved by the current host. It returns 1 if the device is reserved by another host.

2. To query and display the persistent reservation on a device, enter **pcmquerypr -vh /dev/hdisk30**.

Same as Example 1. In addition, it displays the persistent reservation key.

3. To query and display which type of persistent reservation is on a device, enter **pcmquerypr -Vh /dev/hdisk#**.

The following output indicates there is SCSI-2 reserve on the device:

```
# pcmquerypr -Vh /dev/hdisk27
connection type: fscsi3
open dev: /dev/hdisk27
```

Attempt to read reservation key...

```
*> ioctl(PR_READ) error; errno = 5 (I/O error)
*> status_validity=0x1, scsi_bus_status=0x2
```

Attempt to read reservation key...

```
*> ioctl(PR_READ) error; errno = 5 (I/O error)
*> status_validity=0x1, scsi_bus_status=0x18
```

Attempt to read reservation key...

```
*> ioctl(PR_READ) error; errno = 5 (I/O error)
*> status_validity=0x1, scsi_bus_status=0x18
```

Attempt to read reservation key...

```
*> ioctl(PR_READ) error; errno = 5 (I/O error)
*> status_validity=0x1, scsi_bus_status=0x18
```

The following output indicates that there is SCSI-3 reserve on the device:

```
# pcmquerypr -Vh /dev/hdisk43
connection type: fscsi0
open dev: /dev/hdisk43
```

Attempt to read reservation key...

```
*> ioctl(PR_READ) error; errno = 5 (I/O error)
*> status_validity=0x1, scsi_bus_status=0x2
```

Attempt to read reservation key...

Attempt to read registration keys...

Read Keys parameter

Generation : 12

Additional Length: 32

Key0 : 0x3236303232344446

Key1 : 0x3236303232344446

Key2 : 0x3236303232344446

Key3 : 0x3236303232344446

resrvpolicy= no_reserve

Reserve Key provided by current host = none (hex)0924ffff

Reserve Key on the device: 0x3236303232344446

Reservation key type: 0x6

Device is reserved by SDD device.

4. To release the persistent reservation if the device is reserved by the current host, enter **pcmquerypr -rh /dev/hdisk30**.

This command releases the persistent reserve if the device is reserved by the current host. It returns 0 if the command succeeds or the device is not reserved. It returns 2 if the command fails.

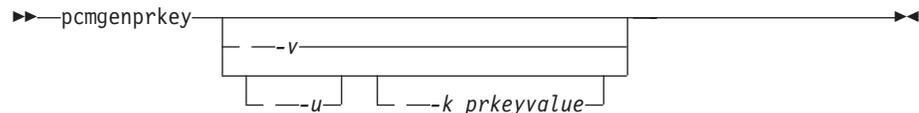
- To reset any persistent reserve and clear all reservation key registrations, enter **pcmquerypr -ch /dev/hdisk30**.
This command resets any persistent reserve and clears all reservation key registrations on a device. It returns 0 if the command succeeds, or 2 if the command fails.
- To remove the persistent reservation if the device is reserved by another host, enter **pcmquerypr -ph /dev/hdisk30**.
This command removes an existing registration and persistent reserve from another host. It returns 0 if the command succeeds or if the device is not persistent reserved. It returns 2 if the command fails.

pcmggenprkey

Description

The **pcmggenprkey** command can be used to set or clear the PR_key_value ODM attribute for all SDDPCM MPIO devices. It also can be used to query and display the reservation policy of all SDDPCM MPIO devices and the persistent reserve key, if those devices have a PR key.

Syntax



Examples

- To set the persistent reserve key to all SDDPCM MPIO devices with a provided key value, issue **pcmggenprkey -u -k 0x1234567890abcdef**. This creates a customized PR_key_value attribute with the provided key value for all SDDPCM MPIO devices, except the devices that already have the same customized PR key attribute. The provided key must contain either a decimal integer or a hexadecimal integer.
- To clear the PR_key_value attribute from all SDDPCM MPIO devices, issue **pcmggenprkey -u -k none**.
- To update the customized PR_key_value attribute with the HACMP-provided Preserve key or the output string from the **uname** command for all the SDDPCM MPIO devices, issue **pcmggenprkey -u**. When the **-u** option is used without the **-k** option, this command searches for the HACMP-provided Preservekey attribute and uses that value as the PR key if that attribute is available; otherwise, it uses the output string from the **uname** command as the PR key.
- To display the reserve_policy, the PR_key_value attribute, and the persistent reserve key attribute of all the SDDPCM devices, issue **pcmggenprkey -v**. If the MPIO device does not have a persistent reserve key, a value of *none* is displayed.

sddpcm_get_config

Description

The **sddpcm_get_config** command can be used to display information about MPIO-based DS4000 or DS5000 subsystems and the hdisks associated with them. Specifically, it displays information about the frame (subsystem), including the frame's assigned name and worldwide name, and a list of hdisks (only those currently in the *Available* state) that are

associated with that subsystem, including the hdisk name, LUN number, current ownership, preferred path, and the user-assigned label for that volume.

Attention: To protect the configuration database, the `sddpcm_get_config` command is not interruptible, because stopping this command before it completes could result in a corrupted database.

Syntax



Flags

-l hdiskN

List information for the subsystem which includes hdiskN.

-A

List information for all attached subsystems.

-V

List additional information, largely of limited value, including the MPIO SDDPCM internal frame number, number of controllers, partition number, and partition count.

Example

Display information about the frame of which hdisk4 is a member.

```
sddpcm_get_config -l hdisk4
```

The system displays a message similar to the following:

```
Storage Subsystem Name = 'SDD_DS5300'
  hdisk      LUN #  Ownership      User Label
  hdisk34    0      B (preferred)  ocotillop2_1
  hdisk35    1      A (preferred)  ocotillop2_2
  hdisk36    2      B (preferred)  ocotillop2_3
  hdisk37    3      A (preferred)  ocotillop2_4
```

Using SDDPCM pcmpath commands

SDDPCM provides commands that you can use to display the status of adapters that are used to access managed devices, to display the status of devices that the device driver manages, or to map supported storage MPIO devices or paths to a supported storage device location. You can also set individual path conditions either to online or offline, set all paths that are connected to an adapter either to online or offline, or set all paths that are connected to a supported storage device port or ports to online or offline. This section includes descriptions of these commands. Table 14 provides an alphabetical list of these commands, a brief description, and where to go in this chapter for more information.

Table 14. Commands

Command	Description	Page
<code>pcmpath clear device count</code>	Dynamically clears the error count or error/select counts to zero.	“pcmpath clear device count” on page 147
<code>pcmpath disable ports</code>	Places paths connected to certain ports offline.	“pcmpath disable ports” on page 147

Table 14. Commands (continued)

Command	Description	Page
pcmpath enable ports	Places paths connected to certain ports online.	"pcmpath disable ports" on page 147
pcmpath open device path	Opens an INVALID path.	"pcmpath open device path" on page 150
pcmpath query adapter	Displays information about adapters.	"pcmpath query adapter" on page 152
pcmpath query adaptstats	Displays performance information for all FCS adapters that are attached to SDDPCM devices.	"pcmpath query adaptstats" on page 153
pcmpath query device	Displays information about devices.	"pcmpath query device" on page 154
pcmpath query devstats	Displays performance information for a single SDDPCM device or all SDDPCM devices.	"pcmpath query devstats" on page 158
pcmpath query essmap	Displays each device, path, location, and attributes.	"pcmpath query essmap" on page 160
pcmpath query port	Displays information about a single target port or all target ports that are attached to SDDPCM-configured MPIO devices.	"pcmpath query port" on page 161
pcmpath query portmap	Displays the status of the logic paths that are managed by SDDPCM between the host and the storage ports.	"pcmpath query portmap" on page 162
pcmpath query portstats	Displays performance information about a single target port or all target ports that are attached to SDDPCM-configured MPIO devices.	"pcmpath query portstats" on page 163
pcmpath query session	Displays the session of the Open HyperSwap devices configured on the host.	"pcmpath query session" on page 165
pcmpath query version	Displays the version of the currently installed SDDPCM.	"pcmpath query version" on page 166
pcmpath query wwpn	Displays the world wide port name (WWPN) for all fibre-channel adapters.	"pcmpath query wwpn" on page 166
pcmpath set adapter	Sets all device paths that are attached to an adapter to online or offline.	"pcmpath set adapter" on page 166
pcmpath set device path	Sets the path of a device to online or offline.	"pcmpath set device path" on page 171
pcmpath set device algorithm	Set all or some of supported storage MPIO device path selection algorithm	"pcmpath set device algorithm" on page 167
pcmpath set device hc_interval	Set all or some of supported storage MPIO device health check time interval	"pcmpath set device hc_interval" on page 168
pcmpath set device hc_mode	Set all or some of supported storage MPIO device health check mode	"pcmpath set device hc_mode" on page 169
pcmpath set device cntlhc_interval	Set all or some of supported active/passive MPIO device controller health check time interval	"pcmpath set device cntlhc_interval" on page 170
pcmpath set device cntlch_delay	Set all or some of supported active/passive MPIO device controller health check delay_time	"pcmpath set device cntlch_delay" on page 170
pcmpath set device svc_sb_ttl	Set the time to live attribute for hdisks in the stretched cluster environment of SVC	"pcmpath set device svc_sb_ttl" on page 172

Parameters

connection

The connection code must be in one of the following formats:

- Single port = R1-Bx-Hy-Zz
- All ports on card = R1-Bx-Hy
- All ports on bay = R1-Bx

Use the output of the **pcmpath query essmap** command to determine the connection code.

essid

The supported storage device serial number, given by the output of **pcmpath query portmap** command.

Examples

If you enter the **pcmpath disable ports R1-B1-H3 ess 12028** command and then enter the **pcmpath query device** command, the following output is displayed:

```
DEV#: 3 DEVICE NAME: hdisk3 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20712028
=====
Path#      Adapter/Path Name    State  Mode      Select  Errors
  0         fscsi0/path0        CLOSE  OFFLINE    6        0
  1         fscsi0/path1        CLOSE  NORMAL     9        0
  2         fscsil/path2        CLOSE  OFFLINE   11        0
  3         fscsil/path3        CLOSE  NORMAL     9        0

DEV#: 4 DEVICE NAME: hdisk4 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20712028
=====
Path#      Adapter/Path Name    State  Mode      Select  Errors
  0         fscsi0/path0        CLOSE  OFFLINE  8702    0
  1         fscsi0/path1        CLOSE  NORMAL   8800    0
  2         fscsil/path2        CLOSE  OFFLINE  8816    0
  3         fscsil/path3        CLOSE  NORMAL   8644    0

DEV#: 5 DEVICE NAME: hdisk5 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20912028
=====
Path#      Adapter/Path Name    State  Mode      Select  Errors
  0         fscsi0/path0        CLOSE  OFFLINE  8917    0
  1         fscsi0/path1        CLOSE  NORMAL   8919    0
  2         fscsil/path2        CLOSE  OFFLINE  9008    0
  3         fscsil/path3        CLOSE  NORMAL   8944    0

DEV#: 6 DEVICE NAME: hdisk6 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20B12028
=====
Path#      Adapter/Path Name    State  Mode      Select  Errors
  0         fscsi0/path0        CLOSE  OFFLINE  9044    0
  1         fscsi0/path1        CLOSE  NORMAL   9084    0
  2         fscsil/path2        CLOSE  OFFLINE  9048    0
  3         fscsil/path3        CLOSE  NORMAL   8851    0

DEV#: 7 DEVICE NAME: hdisk7 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20F12028
=====
Path#      Adapter/Path Name    State  Mode      Select  Errors
  0         fscsi0/path0        CLOSE  OFFLINE  9089    0
  1         fscsi0/path1        CLOSE  NORMAL   9238    0
  2         fscsil/path2        CLOSE  OFFLINE  9132    0
```

```

3          fscsil/path3          CLOSE    NORMAL    9294      0
DEV#:  8  DEVICE NAME: hdisk8  TYPE: 2105E20  ALGORITHM: Load Balance
SERIAL: 21012028
=====
Path#      Adapter/Path Name      State   Mode      Select    Errors
0          fscsi0/path0              CLOSE   OFFLINE   9059      0
1          fscsi0/path1              CLOSE   NORMAL    9121      0
2          fscsil/path2              CLOSE   OFFLINE   9143      0
3          fscsil/path3              CLOSE   NORMAL    9073      0

```

pcmpath enable ports

The **pcmpath enable ports** command sets MPIO device paths online for the specified supported storage device location code.

Syntax

```

▶▶—pcmpath enable ports—connection—ess essid—————▶▶

```

Parameters

connection

The connection code must be in one of the following formats:

- Single port = R1-Bx-Hy-Zz
- All ports on card = R1-Bx-Hy
- All ports on bay = R1-Bx

Use the output of the **pcmpath query essmap** command to determine the connection code.

essid

The supported storage device serial number, given by the output of **pcmpath query portmap** command.

Examples

If you enter the **pcmpath enable ports R1-B1-H3 ess 12028** command and then enter the **pcmpath query device** command, the following output is displayed:

```

DEV#:  3  DEVICE NAME: hdisk3  TYPE: 2105E20  ALGORITHM: Load Balance
SERIAL: 20112028
=====
Path#      Adapter/Path Name      State   Mode      Select    Errors
0          fscsi0/path0              CLOSE   NORMAL     6         0
1          fscsi0/path1              CLOSE   NORMAL     9         0
2          fscsil/path2              CLOSE   NORMAL    11         0
3          fscsil/path3              CLOSE   NORMAL     9         0

DEV#:  4  DEVICE NAME: hdisk4  TYPE: 2105E20  ALGORITHM: Load Balance
SERIAL: 20712028
=====
Path#      Adapter/Path Name      State   Mode      Select    Errors
0          fscsi0/path0              CLOSE   NORMAL    8702      0
1          fscsi0/path1              CLOSE   NORMAL    8800      0
2          fscsil/path2              CLOSE   NORMAL    8816      0
3          fscsil/path3              CLOSE   NORMAL    8644      0

```

```

DEV#:  5  DEVICE NAME: hdisk5  TYPE: 2105E20  ALGORITHM: Load Balance

```

SERIAL: 20912028

```
=====
```

Path#	Adapter/Path Name	State	Mode	Select	Errors
0	fscsi0/path0	CLOSE	NORMAL	8917	0
1	fscsi0/path1	CLOSE	NORMAL	8919	0
2	fscsi1/path2	CLOSE	NORMAL	9008	0
3	fscsi1/path3	CLOSE	NORMAL	8944	0

DEV#: 6 DEVICE NAME: hdisk6 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20B12028

```
=====
```

Path#	Adapter/Path Name	State	Mode	Select	Errors
0	fscsi0/path0	CLOSE	NORMAL	9044	0
1	fscsi0/path1	CLOSE	NORMAL	9084	0
2	fscsi1/path2	CLOSE	NORMAL	9048	0
3	fscsi1/path3	CLOSE	NORMAL	8851	0

DEV#: 7 DEVICE NAME: hdisk7 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20F12028

```
=====
```

Path#	Adapter/Path Name	State	Mode	Select	Errors
0	fscsi0/path0	CLOSE	NORMAL	9089	0
1	fscsi0/path1	CLOSE	NORMAL	9238	0
2	fscsi1/path2	CLOSE	NORMAL	9132	0
3	fscsi1/path3	CLOSE	NORMAL	9294	0

DEV#: 8 DEVICE NAME: hdisk8 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 21012028

```
=====
```

Path#	Adapter/Path Name	State	Mode	Select	Errors
0	fscsi0/path0	CLOSE	NORMAL	9059	0
1	fscsi0/path1	CLOSE	NORMAL	9121	0
2	fscsi1/path2	CLOSE	NORMAL	9143	0
3	fscsi1/path3	CLOSE	NORMAL	9073	0

pcmpath open device path

The **pcmpath open device path** command dynamically opens a path that is in the invalid state. You can use this command to open an invalid path even when I/O is actively running on the devices.

Syntax

```
▶▶—pcmpath open device—device number—path—path number—▶▶
```

Parameters

device number

The logical device number of this hdisk, as displayed by the **pcmpath query device** command.

path number

The path ID that you want to change, as displayed under “Path Name” by the **pcmpath query device** command.

Examples

If you enter the **pcmpath query device 23** command, the following output is displayed:

DEV#: 23 DEVICE NAME: hdisk23 TYPE: 2105E20 ALGORITHM: Load Balance
 SERIAL: 20112028

```
=====
```

Path#	Adapter/Path Name	State	Mode	Select	Errors
0	fscsi1/path0	OPEN	NORMAL	557	0
1	fscsi1/path1	OPEN	NORMAL	568	0
2	fscsi0/path2	INVALID	NORMAL	0	0
3	fscsi0/path3	INVALID	NORMAL	0	0

Note that the current state of path 2 and path 3 is INVALID, which means that open path 2 and path 3 failed.

If the root cause of the path 2 open failure is fixed and you enter the **pcmpath open device 23 path 2** command, the following output is displayed:

Success: device 23 path 2 opened

DEV#: 23 DEVICE NAME: hdisk23 TYPE: 2105E20 ALGORITHM: Load Balance
 SERIAL: 20112028

```
=====
```

Path#	Adapter/Path Name	State	Mode	Select	Errors
0	fscsi1/path0	OPEN	NORMAL	557	0
1	fscsi1/path1	OPEN	NORMAL	568	0
2	fscsi0/path2	OPEN	NORMAL	0	0
3	fscsi0/path3	INVALID	NORMAL	0	0

After issuing the **pcmpath open device 23 path 2** command, the state of path 2 becomes OPEN.

The terms used in the output are defined as follows:

Dev# The logical device number of this hdisk.

Device name
 The name of this device.

Type The device product ID from inquiry data.

Algorithm
 The current path selection algorithm for the device. The algorithm selected is one of the following types: load balancing, load balancing port, round robin, or failover.

Serial The LUN for this device.

Path# The path index displayed by the **pcmpath query device** command.

Adapter
 The name of the adapter to which the path is attached.

Path Name
 The name of the path. The number displayed as part of the name is the path ID of this path that is used by the **pcmpath open device path** and **pcmpath set device path** commands.

State The condition of each path of the named device:

Open Path is in use.

Close Path is not being used.

Close_Failed
 Path is broken and is not being used.

Failed Path is opened, but no longer functional because of error.

Invalid

The path failed to open.

Mode The mode of the named path, which is either *Normal* or *Offline*.

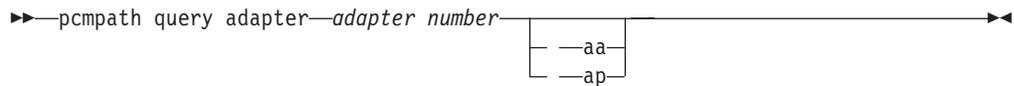
Select The number of times this path was selected for I/O.

Errors The number of I/O errors that occurred on this path.

pcmpath query adapter

The **pcmpath query adapter** command displays information about a single adapter or all adapters that are attached to SDDPCM-configured MPIO devices.

Syntax



Parameters

adapter number

The index number of the adapter for which you want information displayed. If you do not enter an adapter index number, information about all adapters is displayed.

aa The adapter of active/active storage controller devices.

ap The adapter of active/passive storage controller devices.

Examples

If you enter the **pcmpath query adapter** command and your system has both Dual Active or Active/Asymmetric (for example, ESS) and Active/Passive (for example, DS4800) devices, the following output is displayed:

Total Dual Active and Active/Asymmetric Adapters : 2

Adpt#	Name	State	Mode	Select	Errors	Paths	Active
0	fscsi2	NORMAL	ACTIVE	920506	0	80	38
1	fscsi0	NORMAL	ACTIVE	921100	0	80	38

Total Active/Passive Adapters : 2

Adpt#	Name	State	Mode	Select	Errors	Paths	Active
0	fscsi0	NORMAL	ACTIVE	0	0	30	0
1	fscsi1	NORMAL	ACTIVE	0	0	30	0

If you enter the **pcmpath query adapter** command on a host with RSSM LUNs, the following output is displayed:

Total Dual Active and Active/Asymmetric Adapters : 1

Adpt#	Name	State	Mode	Select	Errors	Paths	Active
0	sas1	NORMAL	ACTIVE	1680197	0	32	32

The terms used in the output are defined as follows:

Adpt #

The index number of the adapter.

Name The name of the adapter.

State The condition of the named adapter. It can be either:

Normal

Adapter is in use.

Degraded

One or more opened paths are not functioning.

Failed All opened paths that are attached to this adapter are not functioning.

Mode The mode of the named adapter, which is either *Active* or *Offline*.

Select The number of times this adapter was selected for I/O.

Errors The number of errors that occurred on all paths that are attached to this adapter.

Paths The number of paths that are attached to this adapter.

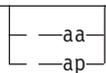
Active The number of functional paths that are attached to this adapter. The number of functional paths is equal to the number of opened paths attached to this adapter minus any that are identified as failed or disabled (offline).

pcmpath query adaptstats

The **pcmpath query adaptstats** command displays information about a single or all fibre-channel adapters that are attached to SDDPCM-configured MPIO devices. If you do not enter an adapter number, information about all adapters is displayed.

Syntax

```
▶▶—pcmpath query adaptstats—adapter number—▶▶
```



Parameters

adapter number

The index number of the adapter for which you want information displayed. If you do not enter an adapter index number, information about all adapters is displayed.

aa The adapter of active/active storage controller devices.

ap The adapter of active/passive storage controller devices.

Examples

If you enter the **pcmpath query adaptstats** command and your system only has Active/Passive devices (for example, DS4800), the following output is displayed:

```
Total Active/Passive Adapters : 2
Adapter #: 0
=====
                Total Read  Total Write  Active Read  Active Write  Maximum
I/O:                0          0           0           0           0
SECTOR:              0          0           0           0           0
Adapter #: 1
=====
                Total Read  Total Write  Active Read  Active Write  Maximum
I/O:                0          0           0           0           0
```

SECTOR: 0 0 0 0 0

/*-----*/

The terms used in the output are defined as follows:

Total Read

- I/O: total number of completed read requests
- SECTOR: total number of sectors that have been read

Total Write

- I/O: total number of completed write requests
- SECTOR: total number of sectors that have been written

Active Read

- I/O: total number of read requests in process
- SECTOR: total number of sectors to read in process

Active Write

- I/O: total number of write requests in process
- SECTOR: total number of sectors to write in process

Maximum

- I/O: the maximum number of queued I/O requests
- SECTOR: the maximum number of queued sectors to Read or Write

pcmpath query device

The **pcmpath query device** command displays information about a single SDDPCM MPIIO device, a set of SDDPCM MPIIO devices between device M and N, a specific model of SDDPCM MPIIO devices, or all SDDPCM MPIIO devices. If two device numbers are entered, the command displays device information for a set of devices between these two device numbers. The option to specify a model cannot be used when you query with device numbers. If you do not enter a device number, information about all devices is displayed. If a device number is entered, the command displays the device information about the hdisk that is associated with this number.

Beginning with SDDPCM 2.1.0.7, a new option, **-d**, is provided to display devices of a particular device model. The valid device models are:

- 1750 - All 1750 models (DS6000)
- 2105 - All 2105 models (ESS)
- 2107 - All 2107 models (DS8000)
- 2145 - All 2145 models (SAN Volume Controller)

Beginning with SDDPCM 2.4.0.0, the following device models are added for option **-d**:

- 1722 - All 1722 devices (DS4300)
- 1724 - All 1724 devices (DS4100)
- 1742 - All 1742 devices (DS4400 and DS4500)
- 1814 - All 1814 devices (DS4200 and DS4700)
- 1815 - All 1815 devices (DS4800)
- 1818 - All 1818 devices (DS5100 and DS5300)

Beginning with SDDPCM 2.6.0.0, the following device model is added for option -d:

- 1820 - All 1820 devices (RSSM)

Beginning with SDDPCM 2.1.0.7, a new feature is available to verify SAN configurations in a controller environment, such as DS6000 or SAN Volume Controller. A nonpreferred path will be marked with "*".

For example, if you configure 4 paths per MPIO hdisk device and you want to configure an equal distribution between preferred controller and nonpreferred controller of DS6000 devices, you can configure your environment with 2 paths connected to a preferred controller and 2 paths connected to a nonpreferred controller. This feature will help you verify your configuration by indicating which path is connected to the nonpreferred controller before the device starts operation.

Note:

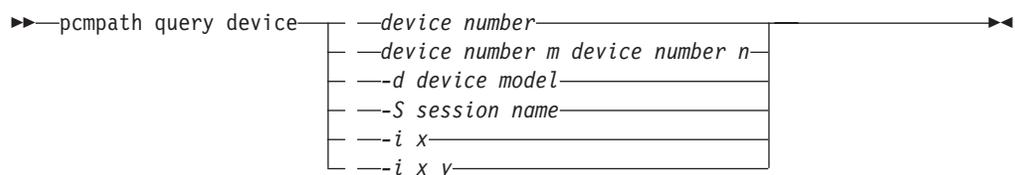
1. You must open SDDPCM MPIO devices of DS6000, and SAN Volume Controller, once before the '*' mark on nonpreferred paths can be displayed in the following three commands:
 - **pcmpath query device**
 - **pcmpath query essmap**
 - **pcmpath query portmap**
2. You must open SDDPCM MPIO RSSM devices before the '*' mark on nonpreferred paths can be displayed in the **pcmpath query device** command. For DS4000 and DS5000 SDDPCM MPIO devices, only opened passive paths are displayed with the '*' mark in the **pcmpath query device** command.

Beginning with SDDPCM 2.1.3.0, two new options are added to the **device query** command. The first option lets you specify two numbers to query a set of devices; the second option *-i x y* lets you repeat the query command every *x* seconds for *y* times.

Beginning with SDDPCM 3.0.0.0, a new **session name** option is added to the **device query** command that allows you to query a session name to show the set of devices in that session.

The **pcmpath query device** commands display only supported storage MPIO devices that are configured with the SDDPCM module. Any AIX internal disks or non-SDDPCM-configured MPIO devices are not displayed.

Syntax



Parameters

device number

The device number refers to the logical device number of the hdisk.

device number_m device_number_n

Use the *device_number_m device_number_n* option to provide a range of device index numbers.

device model

Displays devices of a particular device model. The valid device models are:

- 1722 - All 1722 devices (DS4300)
- 1724 - All 1724 devices (DS4100)
- 1742 - All 1742 devices (DS4400 and DS4500)
- 1750 - All 1750 models (DS6000)
- 1814 - All 1814 devices (DS4200 and DS4700)
- 1815 - All 1815 devices (DS4800)
- 1818 - All 1818 devices (DS5100 and DS5300)
- 2105 - All 2105 models (ESS)
- 2107 - All 2107 models (DS8000)
- 2145 - All 2145 models (SAN Volume Controller)
- 1820 - All 1820 devices (RSSM)

session name

Displays the set of devices in the specified session.

- i* Repeats the command every *x* seconds for *y* times. If you do not specify *y*, the command repeats indefinitely every *x* seconds.

Examples

If you enter the **pcmpath query device 65 66** command, the following output is displayed:

For the supported storage devices:

```
DEV#: 65 DEVICE NAME: hdisk65 TYPE: 1814 ALGORITHM: Load Balance
SERIAL: 600A0B800011796C0000B68148283C9C
=====
Path#      Adapter/Path Name    State   Mode     Select  Errors
  0*       fscsi0/path0        OPEN   NORMAL   42      0
  1*       fscsi1/path2        OPEN   NORMAL   44      0
  2        fscsi0/path1        OPEN   NORMAL   54      0
  3        fscsi1/path3        OPEN   NORMAL   52      0

DEV#: 66 DEVICE NAME: hdisk66 TYPE: 1814 ALGORITHM: Load Balance
SERIAL: 600A0B80001179440000C1FE48283A08
=====
Path#      Adapter/Path Name    State   Mode     Select  Errors
  0        fscsi0/path0        CLOSE  NORMAL   44      1
  1        fscsi1/path2        CLOSE  NORMAL   57      0
  2        fscsi0/path1        CLOSE  NORMAL   0       0
  3        fscsi1/path3        CLOSE  NORMAL   1       0
```

If you enter the **pcmpath query device 4 5** command, the following output is displayed:

```
DEV#: 4 DEVICE NAME: hdisk4 TYPE: 1820N00 ALGORITHM: Load Balance
SERIAL: 6005076B07409F7F4B54DCE9000000B9
=====
Path#      Adapter/Path Name    State   Mode     Select  Errors
  0        sas1/path0          OPEN   NORMAL  1748651  0
  1*       sas1/path1          OPEN   NORMAL   364     0

DEV#: 5 DEVICE NAME: hdisk5 TYPE: 1820N00 ALGORITHM: Load Balance
```

SERIAL: 6005076B07409F7F4B54DCED000000BA

```
=====
Path#      Adapter/Path Name    State   Mode     Select   Errors
  0         sas1/path0            OPEN    NORMAL   1748651    0
  1*        sas1/path1            OPEN    NORMAL    364        0
```

If hdisk2 is an Open HyperSwap device that has never been opened, and you enter **pcmpath query device 2**, the following output is displayed:

```
DEV#: 2 DEVICE NAME: hdisk2 TYPE: 2107900 ALGORITHM: Load Balance
SESSION NAME: session10S
DIRECTION: H1->H2
```

```
=====
PRIMARY SERIAL: 10000000F00
```

```
SECONDARY SERIAL: 20000000F80
```

```
-----
Path#  Adapter/Path Name    State   Mode     Select   Errors
  0     fscsi0/path0         CLOSE   NORMAL    0         0
  1     fscsi0/path1         CLOSE   NORMAL    0         0
  2     fscsi0/path2         CLOSE   NORMAL    0         0
  3     fscsi0/path3         CLOSE   NORMAL    0         0
  4     fscsil/path4         CLOSE   NORMAL    0         0
  5     fscsil/path5         CLOSE   NORMAL    0         0
  6     fscsil/path6         CLOSE   NORMAL    0         0
  7     fscsil/path7         CLOSE   NORMAL    0         0
```

If hdisk2 is an Open HyperSwap device that is being opened, and you enter **pcmpath query device 2**, the following output is displayed.

Note: In the following example, hdisk2 is created on two physical devices. At any time, only one of the two physical devices is active for I/Os. The asterisk (*) indicates the active device to which the current I/Os are being sent.

```
DEV#: 2 DEVICE NAME: hdisk2 TYPE: 2107900 ALGORITHM: Load Balance
SESSION NAME: session1
OS DIRECTION: H1->H2
```

```
=====
PRIMARY SERIAL: 10000000F00 *
```

```
-----
Path#  Adapter/Path Name    State   Mode     Select   Errors
  0     fscsi0/path0         OPEN    NORMAL    8         0
  1     fscsi0/path2         OPEN    NORMAL    9         0
  2     fscsil/path4         OPEN    NORMAL    9         0
  3     fscsil/path5         OPEN    NORMAL    9         0
```

```
SECONDARY SERIAL: 20000000F80
```

```
-----
Path#  Adapter/Path Name    State   Mode     Select   Errors
  4     fscsi0/path1         OPEN    NORMAL    0         0
  5     fscsi0/path3         OPEN    NORMAL    0         0
  6     fscsil/path6         OPEN    NORMAL    0         0
  7     fscsil/path7         OPEN    NORMAL    0         0
```

The terms used in the output are defined as follows:

Dev# The logical device number of this hdisk.

Name The logical name of this device.

Type The device product ID from inquiry data.

Algorithm

The current path selection algorithm selected for the device. The algorithm selected is one of the following: load balancing, load balancing port, round robin, or failover.

Session name

The name of the Tivoli Productivity Center for Replication session in which this Open HyperSwap device is contained.

Serial The LUN for this device.

OS direction

The current Open HyperSwap direction. H1->H2 shows that the I/O is active on the H1 site. H1<-H2 shows that the I/O is active on the H2 site.

Primary serial

The serial number of the volume on the H1 site.

Secondary serial

The serial number of the volume on the H2 site.

Path# The path index displayed by device query command.

Adapter

The name of the adapter to which the path is attached.

Path Name

The name of the path. The number displayed as part of the name is the path ID that is used by **pcmpath open device path** and **pcmpath set device path** commands.

State The condition of the path attached to the named device:

Open Path is in use.

Close Path is not being used.

Failed Path is no longer being used. It has been removed from service due to errors.

Close_Failed

Path was detected to be broken and failed to open when the device was opened. The path stays in Close_Failed state when the device is closed.

Invalid

The path is failed to open, but the MPIO device is opened.

Mode The mode of the named path. The mode can be either *Normal* or *Offline*.

Select The number of times this path was selected for I/O.

Errors The number of input and output errors that occurred on a path of this device.

pcmpath query devstats

The **pcmpath query devstats** command displays performance information for a single SDDPCM MPIO device or all SDDPCM MPIO devices. If you do not enter a device number, information about all devices is displayed. If a device number is entered, the command displays the device information about the hdisk that is associated with this number.

The terms used in the output are defined as follows:

Total Read

- I/O: total number of completed read requests
- SECTOR: total number of sectors that have been read

Total Write

- I/O: total number of completed write requests
- SECTOR: total number of sectors that have been written

Active Read

- I/O: total number of read requests in process
- SECTOR: total number of sectors to read in process

Active Write

- I/O: total number of write requests in process
- SECTOR: total number of sectors to write in process

Maximum

- I/O: the maximum number of queued I/O requests
- SECTOR: the maximum number of queued sectors to read or write

Transfer size

- <= 512: the number of I/O requests received, whose transfer size is 512 bytes or less
- <= 4k: the number of I/O requests received, whose transfer size is 4 KB or less (where KB equals 1024 bytes)
- <= 16K: the number of I/O requests received, whose transfer size is 16 KB or less (where KB equals 1024 bytes)
- <= 64K: the number of I/O requests received, whose transfer size is 64 KB or less (where KB equals 1024 bytes)
- > 64K: the number of I/O requests received, whose transfer size is greater than 64 KB (where KB equals 1024 bytes)

pcmpath query essmap

The **pcmpath query essmap** command displays data collected from the **fcppcmmmap** program after the system has been restarted, or by the **pcmpath** program if any device configuration change occurs after the system has been restarted. The **fcppcmmmap** program issues SCSI commands directly to the disks to collect the information that is displayed. This command is supported for disk storage systems only.

Syntax

▶▶—pcmpath query essmap—————▶▶

Examples

If you enter the **pcmpath query essmap** command, the following output is displayed:

Disk	Path	P	Location	adapter	LUN	SN	Type	Size	LSS	Vol	Rank	C/A	S ...
hdisk5	path0 *	30-60-01	[FC]	fscsi1	13	AAKA1200	IBM 1750-500	1.1	18	0	0000	01	Y ...

```

hdisk5 path1 30-60-01[FC] fscsi0 13AAAKA1200 IBM 1750-500 1.1 18 0 0000 01 Y ...
hdisk5 path2 * 20-60-01[FC] fscsi0 13AAAKA1200 IBM 1750-500 1.1 18 0 0000 01 Y ...
hdisk5 path3 20-60-01[FC] fscsi1 13AAAKA1200 IBM 1750-500 1.1 18 0 0000 01 Y ...

```

The terms used in the output are defined as follows:

Disk The logical device name assigned by the host.

Path The logical path name of a MPIO device.

P Indicates the logical paths and whether the path is preferred and nonpreferred. "*" indicates that the path is a nonpreferred path.

Location

The physical location code of the host adapter through which the LUN is accessed.

Adapter

The logical adapter name assigned by the host LUN.

LUN SN

The unique serial number for each LUN within the supported storage device.

Type The device and model.

Size The capacity of the configured LUN.

LSS The logical subsystem where the LUN resides. (Beginning with 2.1.3.0, the value displayed is changed from decimal to hexadecimal.)

Vol The volume number within the LSS.

Rank The unique identifier for each RAID array within the supported storage device.

C/A The cluster and adapter accessing the array.

S Indicates that the device is shared by two and more supported storage device ports. Valid values are *yes* or *no*.

Connection

The physical location code of supported storage device adapter through which the LUN is accessed.

Port The supported storage device port through which the LUN is accessed.

RaidMode

The disk RAID mode.

pcmpath query port

The **pcmpath query port** command displays information about a single target port or all target ports that are attached to SDDPCM-configured MPIO devices.

Syntax

▶▶—pcmpath query port—*target port number*————▶▶

Parameters

target port number

Use the target port number option to display information about the target port. If you do not enter a target port number, information about all target ports is displayed.

Examples

If you have 12 active ports and enter the **pcmpath query port** command, the following output is displayed:

Port#	Wwpn	State	Mode	Select	Errors	Paths	Active
0	5005022300ca031e	NORMAL	ACTIVE	0	0	60	0
1	5005022307044750	NORMAL	ACTIVE	0	0	64	0
2	5005022307155750	NORMAL	ACTIVE	0	0	64	0
3	500502230e02ac54	NORMAL	ACTIVE	0	0	64	0
4	500502230e01fe54	NORMAL	ACTIVE	0	0	64	0
5	500502230eb3d54	NORMAL	ACTIVE	0	0	64	0
6	500503380110046b	NORMAL	ACTIVE	0	0	60	0
7	5005033801100d64	NORMAL	ACTIVE	0	0	60	0
8	50050338011003d1	NORMAL	ACTIVE	0	0	60	0
9	5005033801100118	NORMAL	ACTIVE	0	0	60	0
10	500503330323006a	NORMAL	ACTIVE	3167	0	62	4
11	500503330323406a	NORMAL	ACTIVE	3220	1	62	4

The terms used in the output are defined as follows:

Port # The index number of the port.

WWPN

The worldwide port name of the target port.

State The condition of the named target port, which can be one of the following types:

Normal

The target port that is in use.

Degraded

One or more of the opened paths are not functioning.

Failed All opened paths that are attached to this target port are not functioning.

Mode The mode of the named target port, which is either *Active* or *Offline*.

Select The number of times this target port was selected for input or output.

Errors The number of errors that occurred on all paths that are attached to this target port.

Paths The number of paths that are attached to this target port.

Active The number of functional paths that are attached to this target port. The number of functional paths is equal to the number of opened paths that are attached to this target port minus any that are identified as failed or disabled (offline).

Note: This command is supported only with 1750, 2105, 2107, and 2145 device types.

pcmpath query portmap

The **pcmpath query portmap** command displays the status of the logic paths that are managed by SDDPCM between the host and the storage ports. This command is supported for disk storage systems only.

Syntax

►—pcmpath query portmap—◄

Examples

If you enter the **pcmpath query portmap** command, the following output is displayed:

```

ESSID    DISK      BAY-1(B1)      BAY-2(B2)      BAY-3(B3)      BAY-4(B4)
          H1 H2 H3 H4      H1 H2 H3 H4      H1 H2 H3 H4      H1 H2 H3 H4
          ABCD ABCD ABCD ABCD  ABCD ABCD ABCD ABCD  ABCD ABCD ABCD ABCD  ABCD ABCD ABCD ABCD
          BAY-5(B5)      BAY-6(B6)      BAY-7(B7)      BAY-8(B8)
          H1 H2 H3 H4      H1 H2 H3 H4      H1 H2 H3 H4      H1 H2 H3 H4
          ABCD ABCD ABCD ABCD  ABCD ABCD ABCD ABCD  ABCD ABCD ABCD ABCD  ABCD ABCD ABCD ABCD

13AAAKA hdisk5 0--- ---- ---- ----  o--- ---- ---- ----  ---- ---- ---- ----  ---- ---- ---- ----
13AAAKA hdisk6 Y--- ---- ---- ----  y--- ---- ---- ----  ---- ---- ---- ----  ---- ---- ---- ----

Y = online/open          y = (alternate path) online/open
O = online/closed       o = (alternate path) online/closed
N = offline             n = (alternate path) offline
- = path not configured
PD = path down

```

The terms used in the output are defined as follows:

- Y** The port is online and open, meaning that at least one path attached to this port is functional.
- y** Paths connected to this port are nonpreferred paths. The port is online and open, meaning that at least one path attached to this port is functional.
- O** The port is online and closed, meaning that at least one path state and mode is closed and online.
- o** Paths connected to this port are nonpreferred paths. The port is online and closed, meaning that at least one path state and mode is closed and online.
- N** The port is offline, meaning that all paths attached to this port are offline.
- n** Paths connected to this port are nonpreferred paths. The port is offline, meaning that all paths attached to this port are offline.
- The path is not configured.
- PD** The path is down. It is either not functional or has been placed offline.

Note: The following fields apply only 1750 devices and can only be shown after the device is opened once:

- y
- o
- n

The serial number of ESS devices is five digits, whereas the serial number of DS6000 and DS8000 devices is seven digits.

pcmpath query portstats

The **pcmpath query portstats** command displays performance information about a single target port or all target ports that are attached to SDDPCM-configured MPIO devices. If you do not enter a device number, information about all devices is displayed.

Syntax

►—pcmpath query portstats—*target port number*—►

Parameters

target port number

Use the target port number option to display information about the target port. If you do not enter a target port number, information about all target ports is displayed.

Examples

If you have four target ports and enter the **pcmpath query portstats** command, the following output is displayed:

```
Port #: 0
=====
              Total Read  Total Write  Active Read  Active Write  Maximum
I/O:          0           0           0           0           0
SECTOR:       0           0           0           0           0

Port #: 1
=====
              Total Read  Total Write  Active Read  Active Write  Maximum
I/O:          0           0           0           0           0
SECTOR:       0           0           0           0           0

Port #: 2
=====
              Total Read  Total Write  Active Read  Active Write  Maximum
I/O:         3109          58           0           0           2
SECTOR:     28292         4084           0           0        2096

Port #: 3
=====
              Total Read  Total Write  Active Read  Active Write  Maximum
I/O:         3169          51           0           0           2
SECTOR:     24762         6470           0           0        2096
```

The terms used in the output are defined as follows:

Total Read

- I/O: The total number of completed read requests
- SECTOR: The total number of sectors that have been read

Total Write

- I/O: The total number of completed write requests
- SECTOR: The total number of sectors that have been written

Active Read

- I/O: The total number of read requests in process
- SECTOR: The total number of sectors to read in process

Active Write

- I/O: The total number of write requests in process
- SECTOR: The total number of sectors to write in process.

Maximum

- I/O: The maximum number of queued I/O requests
- SECTOR: The maximum number of queued sectors to Read or Write

Note:

1. This command is supported only with 1750, 2105, 2107, and 2145 device types.

- Data that is displayed by this command is collected only when the device's algorithm is set to lbp. For example, if the algorithm from hdisk10 to hdisk20 is set to lbp, the statistical data for each device is saved in the associated ports and displayed here. If none of the device's algorithm is set to lbp, there is no change in the port statistical output.

pcmpath query session

The **pcmpath query session** command displays the session of the Open HyperSwap devices configured on the host.

Syntax

```
▶▶—pcmpath query session—————▶▶
```

Parameters

None

Examples

If you enter the **pcmpath query session** command, the following output is displayed:

```
Total Open Hyperswap Sessions : 1
```

```
SESSION NAME: session1
```

SessId	Host_OS_State	Host_Total_copysets	Disabled	Quies	Resum	SwRes
0	READY	2	0	0	0	0

The terms used in the output are defined as follows:

Session Name

The name of the Tivoli Productivity Center for Replication session in which this Open HyperSwap device is contained.

SessId The session ID.

Host_OS_State

The session state on this host (READY or NOT READY).

Host_Total_copysets

The number of Open HyperSwap devices configured on this host.

Disabled

The number of Open HyperSwap devices disabled for HyperSwap on this host.

Quies The number of Open HyperSwap devices quiesced or in the process of being quiesced.

Resum

The number of Open HyperSwap devices in the resume process.

SwRes

Number of Open HyperSwap devices in the swap-and-resume process.

pcmpath query version

The **pcmpath query version** command displays the version of SDDPCM that is installed.

Syntax

```
▶▶—pcmpath query version—▶▶
```

Parameters

None

Examples

If you enter the **pcmpath query version** command, the following output is displayed:

```
[root@abc]> pcmpath query version
IBM SDDPCM Version 2.1.1.0 (devices.sddpcm.52.rte)
```

pcmpath query wwpn

The **pcmpath query wwpn** command displays the World Wide Port Name (WWPN) of the host fibre-channel adapters.

Syntax

```
▶▶—pcmpath query wwpn—▶▶
```

Parameters

None

Examples

If you enter the **pcmpath query wwpn** command, the following output is displayed:

Adapter Name	PortWWN
fscsi0	10000000C925F5B0
fscsi1	10000000C9266FD1

pcmpath set adapter

The **pcmpath set adapter** command sets all device paths attached to an adapter either to online or offline. SDDPCM reserves the last path of a device if the device is opened. This command will fail if there is any device that has the last path attached to this adapter. If all paths of a device are attached to a single fibre-channel adapter that connects to multiple storage ports through a switch, the **pcmpath set adapter 0 offline** command will fail, because the device will lose all its paths.

Attention: If the device is closed, the **pcmpath set adapter offline** command does not have the last path offline protection.

Note: For DS4000 and DS5000 storage devices, this set adapter offline command might cause multiple devices to fail over to an alternate controller if the device path that is attached to this adapter is the last active path. Furthermore, if some devices are owned by controller A and some devices are owned by controller B, the set adapter offline command might cause some devices to fail over from controller A to controller B and some devices to fail over from controller B to controller A. It is possible that the **pcmpath set adapter M offline** command returns a failure message, but some number of devices are failed over successfully and some devices are not. You must issue the **pcmpath query device** command to find out if the command is fully completed or partially completed.

Attention: If the AIX server has more than one DS4000 or DS5000 device model configured, the **pcmpath set adapter N offline** command might cause multiple device models to fail over to the alternate controller. In this case, the **pcmapth set adapter N offline** command is not supported.

Attention: For Open HyperSwap devices, when the device is open, an attempt to set the last active path of the source volume offline by issuing the **pcmpath set adapter** command will be rejected. When the device is closed, it's allowed to set the last path of the source volume offline. However, when the device is opened with all paths to the source volume offline, an unplanned HyperSwap will be triggered.

Syntax

```
▶▶—pcmpath set adapter—adapter number—┌ —online—┐ ┌ —aa—┐ ───────────▶
                                     └ —offline—┘ └ —ap—┘
```

Parameters

adapter number

The index number of the adapter that you want to change.

online

Enables the adapter for service.

offline

Disables the adapter from service.

aa The adapter of active/active storage controller devices.

ap The adapter of active/passive storage controller devices.

Examples

If you enter the **pcmpath set adapter 0 offline ap** command:

- Adapter 0 of the active/passive controller devices changes to *Offline* mode and, if there are some paths in the opened state, its state might change to *failed*.
- All paths of the active/passive controller devices that are attached to adapter 0 change to *Offline* mode and their states change to *Dead*, if they were in the *Open* state.

pcmpath set device algorithm

The **pcmpath set device algorithm** command dynamically changes the path selection algorithm per MPIO device.

Syntax

► `pcmpath set device` `—num1` `—num2` `algorithm` `option` ►

Note: If the device reserve policy is set to `single_path` (SCSI-2 reserve), the device algorithm must be set to `fail_over`. Any attempt to set the algorithm to `round_robin`, `load_balance`, or `load_balance_port` with `single_path` reserve policy will fail.

Parameters

`num1` [`num2`]

- When only `num1` is specified, the command applies to the hdisk specified by `num1`.
- When two device logical numbers are entered, this command applies to all the devices whose logical numbers fit within the range of the two device logical numbers.

`option`

Specifies one of the following path selection algorithms:

- **rr**, where *rr* indicates round robin
- **lb**, where *lb* indicates load balancing
- **fo**, where *fo* indicates failover policy
- **lbp**, where *lbp* indicates load balancing port

Note:

1. You can enter the **pcmpath set device *N* algorithm *rr/fo/lb/lbp*** command to dynamically change to the path selection algorithm associated with SDDPCM MPIO devices that are in either Close or Open state.
2. Beginning with SDDPCM 2.4.0.0, the algorithm *lbp* incorporates I/O statistics from both host adapters and target ports in the path selection algorithm. This new algorithm is applicable only for device models 1750, 2105, 2107, and 2145.

Examples

If you enter **pcmpath set device 2 10 algorithm rr**, the path-selection algorithm of hdisk 2 to hdisk 10 is immediately changed to the round robin algorithm.

You can also use the **chdev** command to change the path selection algorithm of a device: **chdev -l hdiskX -a algorithm=load_balance_port**

pcmpath set device hc_interval

The **pcmpath set device hc_interval** command dynamically changes the health check time interval for SDDPCM MPIO devices or disables the health check function of a device.

Note: Starting with SDDPCM v2123, when you issue this command to change the SDDPCM internal health check function for one or more devices from *enable* to *disable* or from *disable* to *enable*, the SDDPCM server daemon (`pcmsrv`) automatically starts or stops the health check function on the devices.

Syntax

```
►►—pcmpath set device— —num1— [ —num2— ] hc_interval—t—►►
```

Parameters

num1 [*num2*]

- When only *num1* is specified, the command applies to the hdisk specified by *num1*.
- When 2 device logical numbers are entered, this command applies to all the devices whose logical numbers fit within the range of the two device logical numbers.

t The range of supported values for health check interval is 1-3600 seconds. To disable the health check function of a device, set interval time to 0.

Examples

If you enter **pcmpath set device 2 10 hc_interval 30**, the health check time interval of hdisk2 to hdisk10 is immediately changed to 30 seconds.

pcmpath set device hc_mode

The **pcmpath set device hc_mode** command dynamically changes the MPIO device health check mode.

Syntax

```
►►—pcmpath set device— —num1— [ —num2— ] hc_mode—option—►►
```

Parameters

num1 [*num2*]

- When only *num1* is specified, the command applies to the hdisk specified by *num1*.
- When 2 device logical numbers are entered, this command applies to all the devices whose logical numbers fit within the range of the two device logical numbers.

option

Specifies one of the following policies:

- **enabled**, indicates the health check command will be sent to paths that are opened with a normal path mode.
- **failed**, indicates the health check command will be sent to paths that are in failed state.
- **nonactive**, indicates the health check command will be sent to paths that have no active I/O. This includes paths that are opened or in failed state.

Examples

If you enter **pcmpath set device 2 10 hc_mode enabled**, the health check mode of MPIO hdisk2 to hdisk10 is immediately changed to the enabled mode.

pcmpath set device cntlhc_interval

The **pcmpath set device cntlhc_interval** command dynamically changes the Active/Passive MPIIO device controller health check time interval or disables this feature.

Syntax

```
▶▶—pcmpath set device— —num1— [ —num2— ] cntlhc_interval—t—▶▶
```

Parameters

num1 [*num2*]

- When only *num1* is specified, the command applies to the hdisk specified by *num1*.
- When 2 device logical numbers are entered, this command applies to all active/passive devices whose logical numbers fit within the range of the two device logical numbers.

t The range of supported values for controller health check time interval is 0-300 seconds. Setting the value to 0 will disable this feature.

Examples

If you enter **pcmpath set device 2 10 cntlhc_interval 3**, the controller health check time interval of hdisk2 to hdisk10 is immediately changed to 3 seconds, if hdisk2 to hdisk10 are all active/passive devices.

pcmpath set device cntlhc_delay

The **pcmpath set device cntlhc_delay** command dynamically changes the active/passive MPIIO device controller health check delay time or disables this feature.

Syntax

```
▶▶—pcmpath set device— —num1— [ —num2— ] cntlhc_delay—t—▶▶
```

Parameters

num1 [*num2*]

- When only *num1* is specified, the command applies to the hdisk specified by *num1*.
- When 2 device logical numbers are entered, this command applies to all active/passive devices whose logical numbers fit within the range of the two device logical numbers.

t The range of supported values for controller health check time interval is 0-300 seconds. Setting the value to 0 will disable this feature.

Examples

If you enter `pcmpath set device 2 10 cntlhc_delay 30`, the controller health check delay time of `hdisk2` to `hdisk10` is immediately changed to 30 seconds, if `hdisk2` to `hdisk10` are all active/passive devices.

Note:

1. If `cntl_delay_time` is set to '1', it disables the controller health check feature, which is the same as setting it to '0'.
2. If you try to set `cntl_hcheck_int` with a value larger than `cntl_delay_time`, then `cntl_hcheck_int` will be set to the same value as `cntl_delay_time`.
3. If you try to set `cntl_delay_time` with a value smaller than `cntl_hcheck_int`, the command will fail with the INVALID parameter.

pcmpath set device path

The `pcmpath set device path` command sets the path of a device either to *online* or *offline*. You cannot remove the last path to an opened device from service. This prevents a data access failure from occurring.

Attention: If the device is closed, the `pcmpath set adapter offline` command does not have the last path offline protection.

Attention: For Open HyperSwap devices, when the device is open, an attempt to set the last active path of the source volume offline by issuing the `pcmpath set device M path N offline` command will be rejected. When the device is closed, it's allowed to set the last path of the source volume offline. However, when the device is opened with all paths to the source volume offline, an unplanned HyperSwap will be triggered.

Syntax

```
▶▶—pcmpath set device—device number—path—path-ID——online—  
—offline—▶▶
```

Parameters

device number

The logical device number of the hdisk.

path ID

The path ID that you want to change, as displayed under “Path Name” by the `pcmpath query device` command.

online

Enables the path for service.

offline

Disables the path from service.

Examples

If you enter the `pcmpath set device 5 path 0 offline` command, path 0 for device 5 changes to Offline mode.

pcmpath set device svc_sb_ttl

The `pcmpath set device svc_sb_ttl` command dynamically changes the time for which an I/O must remain active after it gets issued.

The SVC storage can encounter a fault scenario called the split-brain condition. In this scenario, the two nodes of the I/O group of an SVC cluster fail to communicate with each other for some time. SVC can recover from this split-brain scenario in a stretched cluster environment in approximately 95 to 120 seconds. During this period, any I/O to the storage is either queued causing a timeout or gets a 'no-device response' error. After this period, one of the nodes performs a target reset and the other node takes over the ownership of the devices.

To provide more time for the storage to establish ownership of the devices, the `svc_sb_ttl` ("SVC time to live") tunable attribute is used. With this attribute set to an appropriate value, SDDPCM detects if all the good paths have been tried at least once and whether the I/O has encountered errors indicating a split-brain condition. In a split-brain scenario, the I/O does not fail immediately, but is retried on all paths again.

Attention: Set the `svc_sb_ttl` attribute **only** for a stretched cluster environment of SVC. In other cases, retain the default attribute value, which is 0.

Syntax

```
▶▶—pcmpath set device— —num1— [ —num2— ]—svc_sb_ttl—t—▶▶
```

Parameters

num1 [*num2*]

- When only *num1* is specified, the command applies to the hdisk specified by *num1*.
- When 2 device logical numbers are entered, this command applies to all the devices whose logical numbers fit within the range of the two device logical numbers.

t The range of supported values for the **time to live** attribute of SVC is 0-240 seconds.

Examples

If you enter `pcmpath set device 2 10 svc_sb_ttl 180`, the **time to live** attribute of hdisk2 to hdisk10 is immediately changed to 180 seconds.

Note: The value of `svc_sb_ttl` must be more than `rw_timeout` so that SVC has additional time to recover from split-brain scenario.

pcmpath chgprefercntl device

The `pcmpath chgprefercntl device` command identifies the preferred controller path to support NDVM.

Starting with SDDPCM 2.6.4.0, the `pcmpath chgprefercntl device` command is added to the `pcmpath` utility. NDVM is an SVC feature that can be used to associate vdisks with two or more io groups, add or remove access to io groups,

and move vdisk access from one io group to another without disruption. For more information about the SAN Volume Controller, see *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

SDDPCM discovers the preferred controller as exported by the target by using in-band commands. For an optimized path control function, SDDPCM must rediscover the change in preferred controller when a vdisk access is moved from one io group to another.

Syntax

► `pcmpath chgprefercntl device—device number` ◄

Parameters

device number

The logical device number of the hdisk.

Examples

On SVC: hdisk5 has 8 paths. Paths 0-3 are from iogroup 0 and paths 4-7 are from iogroup 1. Currently iogroup 0 has access to the vdisk corresponding to hdisk5. Paths 0 and 2 are the preferred paths.

```
# pcmpath query device 5
DEV#: 5 DEVICE NAME: hdisk5 TYPE: 2145 ALGORITHM: Load Balance
SERIAL: 60050768018182DDC0000000000001BE
=====
Path#      Adapter/Hard Disk      State  Mode   Select  Errors
  0         fscsi0/path0          OPEN   NORMAL  51      0
  1*        fscsi0/path1          OPEN   NORMAL  35      0
  2         fscsi0/path2          OPEN   NORMAL  35      0
  3*        fscsi0/path3          OPEN   NORMAL  35      0
  4*        fscsi0/path4          OPEN   NORMAL   5      0
  5*        fscsi0/path5          OPEN   NORMAL   5      0
  6*        fscsi0/path6          OPEN   NORMAL   5      0
  7*        fscsi0/path7          OPEN   NORMAL   5      0
```

To move access of hdisk5 to iogroup 1, run `svctask movevdisk -iogrp 1 <vdisk_id>`.

Now, if you enter the `pcmpath chgprefercntl device 5` command, then SDDPCM rediscover the preferred controllers. After a few seconds, enter `pcmpath query device 5` and see the change in output.

```
# pcmpath query device 5
DEV#: 5 DEVICE NAME: hdisk5 TYPE: 2145 ALGORITHM: Load Balance
SERIAL: 60050768018182DDC0000000000001BE
=====
Path#      Adapter/Hard Disk      State  Mode   Select  Errors
  0*        fscsi0/path0          OPEN   NORMAL  53      0
  1*        fscsi0/path1          OPEN   NORMAL  35      0
  2*        fscsi0/path2          OPEN   NORMAL  38      0
  3*        fscsi0/path3          OPEN   NORMAL  35      0
  4         fscsi0/path4          OPEN   NORMAL   7      0
  5*        fscsi0/path5          OPEN   NORMAL   7      0
  6         fscsi0/path6          OPEN   NORMAL   7      0
  7*        fscsi0/path7          OPEN   NORMAL   7      0
```

Paths 4 and 6 are now the preferred paths.

Summary of command syntax

SDDPCM supports the following **pcmpath** commands:

- **pcmpath clear device** <n>/(<m> <n>) **count error/all**
- **pcmpath query adapter** [<n> aa | ap]
- **pcmpath query adaptstats** [<n> aa | ap]
- **pcmpath query device** <n>/[<m> <n>] [-d <device_model>] [-S <session name>] [-i <x>/(<x> <y>)]
- **pcmpath query devstats** <n>/[<m> <n>] [-d <device_model>] [-i <x>/(<x> <y>)]
- **pcmpath set adapter n** online | offline aa | ap
- **pcmpath set device M path N** online | offline
- **pcmpath set device** <n1> [n2] **algorithm** <option>
- **pcmpath set device** <n1> [n2] **hc_interval** <t>
- **pcmpath set device** <n1> [n2] **hc_mode** <option>
- **pcmpath set device** <n1> [n2] **cntlhc_interval** <t>
- **pcmpath set device** <n1> [n2] **cntlhc_delay** <t>
- **pcmpath set device** <n1> [n2] **svc_sb_ttl** <t>
- **pcmpath disable port** <location> **ess** <essid>
- **pcmpath enable port** <location> **ess** <essid>
- **pcmpath open device** <m> **path** <n>
- **pcmpath query essmap**
- **pcmpath query port** <target port number>
- **pcmpath query portmap**
- **pcmpath query portstats** <target port number>
- **pcmpath query session**
- **pcmpath query wwpn**
- **pcmpath query version**
- **pcmpath chgprefercntl device**

Note: If the commands are used for a device, the *n* is the number of the device logical name. For example, **pcmpath query devstats 3** queries the device statistics for hdisk3.

If the commands are used for adapter, the *n* is the index of the adapter. For example, **pcmpath query adapter 2** queries the adapter statistics for the third adapter in adapter list order, which can be fscsi5.

Chapter 4. Using the SDD on an HP-UX host system

This chapter provides step-by-step procedures for you to install, configure, remove, and use SDD on a Hewlett-Packard (HP-UX) host system that is attached to supported storage devices.

For updated and additional information that is not included in this manual, see the Readme file on the CD-ROM or go to SDD website:

www.ibm.com/servers/storage/support/software/sdd

Verifying the hardware and software requirements

You must install the following hardware and software components to ensure that SDD installs and operates successfully.

Hardware

The following hardware components are needed:

- One or more of the supported storage devices.
- For ESS devices: at least one SCSI host adapter (two are required for load balancing and failover)

To install SDD and use the input/output (I/O) load balancing and failover features, you need a minimum of two SCSI or fibre-channel adapters.

A host system with a single fibre-channel adapter that connects through a switch to multiple ESS ports is considered to have multiple fibre-channel SDD vpath devices.

For information on the fibre-channel adapters that can be used on your HP-UX host system go to:

www.ibm.com/servers/storage/support

- A SCSI cable to connect each SCSI host adapter to a storage system controller port
- Subsystem LUNs that have been created and confirmed for multipoint access
- A fiber-optic cable to connect each fibre-channel adapter to a supported storage device port

Software

SDD supports certain HP-UX kernel levels.

SDD-supported HP-UX kernel levels

SDD supports the following HP-UX kernel levels:

- HP-UX 11.11 (11i) 32-bit and 64-bit
- HP-UX 11.23 (11i v2) PA_RISC 64-bit and IA 64-bit

Ensure that the patches listed in Table 16 on page 178 or newer patches are installed on your HP-UX host system.

Unsupported environments

SDD does not support the following environments:

- HP-UX 11.0 32-bit kernel
- HP-UX 11.0 64-bit kernel
- A system start from an SDD pseudo device
- A system paging file on an SDD pseudo device
- A host system with both a SCSI and fibre-channel connection to a shared LUN
- Single-path mode during concurrent download of licensed machine code nor during any disk storage system concurrent maintenance that impacts the path attachment, such as a disk storage system-host-bay-adapter replacement
- Single-path configuration for fibre channel
- DS8000 and DS6000 with SCSI connectivity

Understanding how SDD works on an HP-UX host system

SDD resides above the HP SCSI disk driver (sdisk) in the protocol stack. For more information about how SDD works, see “The SDD architecture” on page 2.

Preparing for the SDD installation

Before you install the SDD, you must configure the supported storage devices to your host system and attach the required SCSI or fibre-channel adapters. A maximum of 1200 LUNs can be configured by the SDD. The HP-UX SDD supports up to 32 paths per LUN. Be aware that there is no performance gain with a large number of LUNs that are smaller in size. The larger the configuration, the more system resources are required.

Note: 1200 LUNs are supported only by HP-UX 11.23 (PA-RISC 64 and IA64). 600 LUNs are supported by HP-UX 11.11.

Configuring the disk storage system

Before you install SDD, configure your ESS for single-port or multiport access for each LUN. SDD requires a minimum of two independent paths that share the same logical unit to use the load balancing and path-failover-protection features. **With a single path, failover protection is not provided.**

For information about configuring your ESS, see the *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide*.

Configuring the virtualization products

Before you install SDD, configure your SAN Volume Controller devices for single-port or multiport access for each LUN. SDD requires a minimum of two independent paths that share the same logical unit to use the load balancing and path-failover-protection features. **With a single path, failover protection is not provided.**

For information about configuring your SAN Volume Controller, see the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

Configuring DS4000 and DS5000 family disk systems

Before you install the SDD, follow the instructions in the *IBM System Storage DS4x00 Installation, User's and Maintenance Guide* and the *IBM System Storage*

DS4000 Storage Manager Concepts Guide for configuring LUNs that are attached to the HP-UX host system. The SDD requires a minimum of two independent paths that share the same logical unit to use load balance and path failover features. With a single path, failover protection is not provided.

The DS4000 and DS5000 controllers can be set to two different modes: AVT or non-AVT but can operate only in one mode per storage partition. The controller mode is determined by the host type, which is predefined in the DS4000 or DS5000 Storage Manager profile. Only the AVT mode is supported for the HP-UX host type.

To ensure that a single path failure does not result in controller failover during recovery, configure redundant paths to each controller. Because switching to another controller affects performance, configuring the redundant path to a controller can avoid unnecessary controller failover that is caused by a path failure.

Bypassing the boot disk check

The SDD installation process checks the boot disk and then excludes the bootable disk from the SDD configuration. This process attempts to read the boot record from every device instance of the system, so the process can be lengthy if your configuration is large or if the devices are unable to satisfy the I/O request.

If the system boot disk is not an SDD-supported disk, you can bypass the process that checks the boot disk. To do this, enter the following before the SDD installation:

```
touch /opt/IBMsdd/BYPASS_BOOTDISK_CHECK
```

You might want to bypass boot disk checking for systems with DS4000 or DS5000 devices, because reading the boot record from every device instance of DS4000 or DS5000 can be very time consuming.

Planning for installation

Before you install SDD on your HP-UX host system, you need to understand what kind of software runs on your host. The way you install SDD depends on the kind of software that you are running. Two types of special device files are supported:

- Block device files
- Character device files

There are three possible scenarios for installing SDD. The scenario you choose depends on the kind of software that you have installed:

Table 15 further describes the various installation scenarios and how you should proceed. Table 16 on page 178 lists the patches necessary for proper operation of SDD on HP-UX .

Table 15. SDD installation scenarios

Installation scenario	Description	How to proceed
-----------------------	-------------	----------------

Table 15. SDD installation scenarios (continued)

Scenario 1	<ul style="list-style-type: none"> SDD is not installed. The SDD server for Expert is installed. No software application or DBMS communicates directly to sdisk interface. 	Go to: <ol style="list-style-type: none"> “Determining if the SDD 1.3.1.5 (or later) server for Expert is installed” “Installing SDD” on page 179 “Standard UNIX applications” on page 196
Scenario 2	<ul style="list-style-type: none"> SDD is not installed. The SDD server for Expert is installed. An existing application package or DBMS communicates directly to the sdisk interface. 	Go to: <ol style="list-style-type: none"> “Determining if the SDD 1.3.1.5 (or later) server for Expert is installed” “Installing SDD” on page 179 “Using applications with SDD” on page 196
Scenario 3	<ul style="list-style-type: none"> SDD is installed. The SDD server for Expert is installed. 	Go to <ol style="list-style-type: none"> “Determining if the SDD 1.3.1.5 (or later) server for Expert is installed” “Upgrading the SDD” on page 181

Table 16. Patches necessary for proper operation of SDD on HP-UX

HP-UX	Patch bundles
11.23	March 06, standard patch bundles
11.11	September 05, support plus

See www11.itrc.hp.com/service/home/home.do?admit=109447626+1266865784425+28353475 for patch details and prerequisites for patches.

Determining if the SDD 1.3.1.5 (or later) server for Expert is installed

Note: IBM does not support ESS. For more information, contact the IBM Support team.

If you previously installed the SDD server (the stand-alone version) for IBM TotalStorage Expert V2R1 (ESS Expert) on your HP-UX host system, you must remove this stand-alone version of the SDD server before you proceed with SDD 1.3.1.5 installation. The installation package for SDD 1.3.1.5 includes the SDD server daemon (also referred to as `sddsrv`), which incorporates the functionality of the stand-alone version of the SDD server (for ESS Expert).

To determine if the stand-alone version of the SDD server is installed on your host system, enter:

```
swlist SDDsrv
```

If you previously installed the stand-alone version of the SDD server, the output from the `swlist SDDsrv` command looks similar to this:

Note:

1. The installation package for the stand-alone version of the SDD server (for ESS Expert) is SDDsrvHP*bb_yymmdd*.depot (where *bb* represents 32- or 64-bit, and *yymmdd* represents date of installation package). For ESS Expert V2R1, the stand-alone SDD server installation package is SDDsrvHP32_020115.depot for a 32-bit environment, and SDDsrvHP64_020115.depot for a 64-bit environment.)
2. For instructions on how to remove the stand-alone version of the SDD server (for ESS Expert) from your HP-UX host system, see the IBM SUBSYSTEM DEVICE DRIVER SERVER 1.0.0.0 (sddsrv) readme for IBM TotalStorage Expert V2R1 at the following website:
www.ibm.com/servers/storage/support

For more information about the SDD server daemon, go to “SDD server daemon” on page 191.

Installing SDD

Before you install SDD, make sure that you have root access to your HP-UX host system and that all the required hardware and software is ready.

Installing SDD from CD-ROM

Complete the following steps to install SDD on your HP-UX host system:

1. Make sure that the SDD compact disc (CD) is available.
2. Insert the CD-ROM into your CD-ROM drive.
If CD-ROM directory is not known:
 - a. Run `ioscan -funC disk`.
 - b. Record the CD-ROM or DVD-ROM directory.
3. Mount the CD-ROM drive using the **mount** command. Here are two examples of the **mount** command:

```
mount /dev/dsk/c0t2d0 /cdrom
```

or

```
mount /dev/dsk/c0t2d0 /your_installation_directory
```

where `/cdrom` or `/your_installation_directory` is the name of the directory to which you want to mount the CD-ROM drive.

4. Run the **sam** program.

```
> sam
```
5. Select **Software Management**.
6. Select **Install Software to Local Host**.
7. At this point, the **SD Install - Software Selection** panel is displayed. Almost immediately afterwards, a Specify Source menu is displayed:
 - a. For **Source Depot Type**, select the local CD-ROM.
 - b. For **Source Depot Path**, choose the directory and the IBM*sdd*.depot file.

For 32-bit mode applications, use:

```
/cdrom/hp32bit/IBMsdd.depot
```

or

`/your_installation_directory/hp32bit/IBMsdd.depot`

For 64-bit mode applications, use:

`/cdrom/hp64bit/IBMsdd.depot`

or

`/your_installation_directory/hp64bit/IBMsdd.depot`

c. Click **OK**.

You will see output similar to the following example:

IBMsdd_tag 1.7.0.3 IBMsdd Driver 64-bit Version: 1.7.0.3 Sep-24-2007 16:35

8. Click the **IBMsdd_tag** product.
9. From the Bar menu, click **Actions** → **Mark for Install**.
10. From the Bar menu, click **Actions** → **Install (analysis)**. An Install Analysis panel is displayed, showing the status of **Ready**.
11. Click **OK** to proceed. A Confirmation window opens and states that the installation will begin.
12. Click **Yes** and press **Enter**. The analysis phase starts.
13. After the analysis phase has finished, another Confirmation window opens informing you that the system will be restarted after installation is complete. Click **Yes** and press **Enter**. The installation of IBMsdd will now proceed.
14. An Install window opens, informing you about the progress of the IBMsdd software installation. The window looks similar to the following:

```
Press 'Product Summary' and/or 'Logfile' for more target information.
Target          : XXXXX
Status          : Executing install setup
Percent Complete : 17%
Kbytes Installed : 276 of 1393
Time Left (minutes) : 1
Product Summary  Logfile
Done                                     Help
```

The **Done** option is not available when the installation is in progress. It becomes available after the installation process is complete.

15. Click **Done**.

Note: SDD 1.5.0.4 is changed from a static driver to a dynamic loadable kernel module (DLKM) driver. The system does not restart after SDD is installed.

After the installation is finished, the SDD driver is automatically loaded.

You can use the **datapath query device** command to verify the SDD installation. SDD is successfully installed if the command runs successfully.

Installing SDD from downloaded code

The SDD download website, www-01.ibm.com/support/docview.wss?uid=ssg1S4000053, has several SDD packages for different versions of HP-UX. The SDD packages are in tar archives. After you have downloaded the appropriate SDD package into your chosen directory, you can untar the file using the **tar xvf IBMsdd*.tar** command. The complete the following steps to install SDD:

1. Run the sam program.

- > sam
2. Select **Software Management**.
 3. Select **Install Software to Local Host**.
 4. At this point, the SD Install - Software Selection panel is displayed. Then a Specify Source menu is displayed:
 - a. Select the Local Directory for **Source Depot Type**.
 - b. Select the directory in which you have issued the **tar xvf IBMsdd*.tar** command to untar the file and the IBMsdd.depot file for the **Source Depot Path**. Use the fully-qualified path name for the depot file as shown below.
`/your_installation_directory/IBMsdd.depot`
 - c. Click **OK**. Output similar to the following example is displayed:

Name	Revision	Information	Size(Kb)	Architecture	Category
IBMsdd_tag	-> B.11.23.01	IBMsdd Driver 64-bit <version> <date>	nnnn	HPUX_B.11.23_ia64	
 5. Click the **IBMsdd_tag** product and complete the steps beginning with step 9 on page 180 shown in "Installing SDD from CD-ROM" on page 179.

Upgrading the SDD

Use these sections to upgrade SDD:

Upgrading from SDD 1.6.0.x to SDD 1.6.1.0 or later with concurrent access

The memory management and the installation process have been enhanced to allow installation of the SDD package while the LVM volume groups are active and user applications are running. The concurrent driver upgrade function permits uninterrupted operation when installing SDD.

The installation process:

1. Converts SDD vpath devices to PVLINK devices
2. Unloads and reloads the SDD driver
3. Converts the PVLINK devices back to SDD vpath devices after the new package is installed.

Because the volume groups must be active for the PVLINK conversion process, the following are the limitations:

1. The volume groups must be managed by HP-UX LVM.
2. The MC Service Guard cluster must be halted prior to upgrade. The primary node and the adoptive node or nodes must operate in a single-host environment. The shared volume groups in the adoptive nodes must be exported so that the volumes are not shared; the volume groups can be active in the primary node only. Restore the cluster environment after upgrading SDD.

Performance during upgrade: You should consider the following performance topics while you are upgrading:

- The PVLINK conversion process and the driver reload require additional system resources such as LVM lock, accessing LVM meta data and the kernel memory. With the concurrent I/O, the upgrade process can take longer because the conversion process must wait for the I/O to complete before a link can be removed from PVLINK.
- Reloading the SDD driver can also take longer because of the contention with the kernel memory; the system must wait for a window when the resources become available. The actual time for installation depends on the processor

model, physical memory size, I/O intensity, and configuration size. The larger the SDD configuration or the more concurrent I/O activities, the longer it can take to upgrade. The installation time can also take longer if the devices from the ioscan output are not accessible. If there were a lot of inaccessible devices as the result of fabric reconfiguration, you should attempt to clean up the configuration before upgrading.

- After the upgrade, you should check the VPATH_EVENT for allocation failures in syslog.log, /var/adm/IBMsdd/hd2vp.errlog and vp2hd.errlog. These are the indications that the upper limit of the resources have been reached during the conversion process and that you should take a more conservative approach next time. That is, the concurrent upgrade should be performed during a period when the system load is lighter than the normal operation.
- The installation process also ensures that the current SDD state is not in any degraded state; the recovery process can be lengthy if the upgrade failed due to the hardware errors. Issue the **swjob** command that is indicated at the end of swinstall output to get detailed information about the installation.
- The diagnose message in the package installation and configuration process has been greatly improved to include logs for **cfgvpath**, **vp2hd**, **hd2vp** and the syslog messages. All the SDD related logs have been moved to the /var/adm/IBMsdd directory.

Upgrading from SDD 1.5.0.4 to SDD 1.6.1.0 or later with nonconcurrent access.

Upgrading SDD consists of removing and reinstalling the IBMsdd package. If you are upgrading SDD, go to “Uninstalling SDD” on page 190 and then go to “Installing SDD” on page 179.

Upgrading SDD 1.3.0.2 or earlier to SDD 1.5.0.4 or later

SDD 1.3.0.2 (or earlier) could have created a vpath0 device. vpath0 should not be configured. To upgrade your SDD, you need to ensure that you have no vpath0 devices using the following steps:

1. Remove its /etc/vpathsave.cfg (vpath name reservation file) file.
2. Issue **cfgvpath -c** to reconfigure SDD.

To upgrade from SDD 1.3.0.2 (or earlier) to SDD 1.5.0.4 (or later):

1. Uninstall SDD 1.3.0.2 (or earlier). Because SDD 1.3.0.2 (or earlier) does not have the volume group conversion scripts **hd2vp** and **vp2hd**, the uninstallation will not convert the existing volume groups back to cXtXdX devices.
2. Install SDD 1.5.0.4 (or later) on your host.
3. After installation, check to see that you have:
 - a. The volume group conversion scripts **hd2vp** and **vp2hd** in /opt/IBMdpo/bin
 - b. The /etc/vpathsave.cfg file
4. Use the **vp2hd** script to convert volume groups from volume groups containing SDD vpath devices to volume groups containing cXtXdX devices.
5. When the volume groups have been converted:
 - a. Remove the /etc/vpathsave.cfg file.
 - b. Issue **cfgvpath -c** to reboot.
6. When the host responds, use the **showvpath** command to determine that vpath0 has been removed. If vpath0 no longer exists and the state of all the

SDD vpath devices looks good in the **datapath query device** output, use the **hd2vp** script to change volume groups to volume groups containing SDD vpath devices.

Note: Vpathname vpathN is reserved when it is assigned to a LUN even after the LUN has been removed from the host. The same vpathname, vpathN, will be assigned to the same LUN when it is reconnected to the host.

7. `/etc/vpathsave.cfg` is the file to reserve vpathnames. Improper removal of the file will invalidate existing volume groups. Do not remove the `/etc/vpathsave.cfg` file.

Installing an additional package for DS4000 and DS5000 support

The HP SDD support for DS4000 and DS5000 is enabled through the On-Demand-Device function, which recognizes the new devices as a plug-in to the driver. You can install the On-Demand package before or after you install the traditional SDD package.

Note: DS5000 support is limited to HP-UX 11.23 systems, including both IA64 and PA-RISC architectures.

If the On-Demand package is installed later, run **cfgvpath -r** to recognize and reconfigure the DS4000 or DS5000 device. If the On-Demand package was installed earlier, the installation process automatically recognizes and configures the DS4000 and DS5000 devices.

To install the DS On-Demand-Device package, issue the following command:

For DS4000: **swinstall IBMsdd_ds4k_platform.depot IBMsdd_ds4k_tag**

For DS5000: **swinstall IBMsdd_ds5k_platform.depot IBMsdd_ds5k_tag**

Configuring the SDD

This section provides information necessary to configure the SDD. Use the HP command line interface (CLI) to manage SDD devices.

Maximum number of LUNs

A maximum of 1200 LUNs can be configured by SDD. HP-UX SDD supports up to 32 paths per LUN.

Changing an SDD hardware configuration

When you are adding or removing multiport SCSI devices, you must reconfigure SDD to recognize the new devices. Complete the following steps to reconfigure SDD:

1. Issue the **cfgvpath** command to reconfigure the SDD vpath device by entering:

```
/opt/IBMsdd/bin/cfgvpath
```

2. Restart the system by entering:

```
shutdown -r 0
```

Note: Restarting the system is no longer required for SDD 1.5.0.4 or later.

The **quersn** command can be used to list all disk storage system devices visible to the host. The **quersn** command reads the unique serial number of a disk storage system device (sdisk). To manually exclude devices from the SDD configuration, their serial number information can be included in the `/etc/vpathmanualexcl.cfg` text file. For bootable devices, the **get_root_disks** command generates a file called `/etc/vpathexcl.cfg` to exclude bootable disks from the SDD configuration.

Converting a volume group

SDD provides the following conversion scripts:

hd2vp The **hd2vp** script converts a volume group from supported storage device sdisks into SDD vpath devices.

The syntax for **hd2vp** script is as follows:

```
hd2vp vname
```

vp2hd The **vp2hd** script converts a volume group from SDD vpath devices into supported storage device sdisks. Use the **vp2hd** program when you want to configure your applications back to original supported storage device sdisks.

The syntax for **vp2hd** script is as follows:

```
vp2hd vname
```

hd2vp and **vp2hd** convert a volume group from an sdisk pmlink to and from SDD vpaths devices. The **hd2vp** conversion program is invoked at the system boot time to ensure the conversion of sdisks to SDD vpath devices. Beginning with SDD 1.6.0.12, **vp2hd** is not invoked at the system shutdown.

Dynamic reconfiguration

Dynamic reconfiguration provides a way to automatically detect path configuration changes without requiring a reboot.

1. **cfgvpath -r**:

This operation finds the current hardware configuration and compares it to the SDD vpath device configuration in memory and then identifies a list of differences. It then issues commands to update the SDD vpath device configuration in memory with the current hardware configuration. The commands that **cfgvpath -r** issues to the vpath driver are:

- Add an SDD vpath device.
- Remove an SDD vpath device; this will fail if device is busy.
- Add path to the SDD vpath device.
- Remove path from the SDD vpath device; this will fail deletion of the path if the device is busy, but will set path to DEAD and OFFLINE.

2. **rmvpath** command removes one or more SDD vpath devices.

```
rmvpath -all                # Remove all SDD vpath devices

rmvpath vpath_name          # Remove one SDD vpath device at a time
                             # this will fail if device is busy
```

Dynamically changing the SDD path-selection policy algorithm

SDD 1.4.0.0 (or later) supports path-selection policies that increase the performance of multipath-configured supported storage devices and make path failures transparent to applications. The following path-selection policies are supported:

failover only (fo)

All I/O operations for the device are sent to the same (preferred) path until the path fails because of I/O errors. Then an alternate path is chosen for subsequent I/O operations.

load balancing (lb)

The path to use for an I/O operation is chosen by estimating the load on the adapter to which each path is attached. The load is a function of the number of I/O operations currently in process. If multiple paths have the same load, a path is chosen at random from those paths. Load balancing mode also incorporates failover protection.

Note: The load balancing policy is also known as the optimized policy.

round robin (rr)

The path to use for each I/O operation is chosen at random from those paths that were not used for the last I/O operation. If a device has only two paths, SDD alternates between the two.

The path-selection policy is set at the SDD device level. The default path-selection policy for an SDD device is load balancing. You can change the policy for an SDD device. SDD version 1.4.0.0 (or later) supports dynamic changing of the SDD devices path-selection policy.

Before changing the path-selection policy, determine the active policy for the device. Enter **datapath query device N** where *N* is the device number of the SDD vpath device to show the current active policy for that device.

datapath set device policy command

Use the **datapath set device policy** command to change the SDD path-selection policy dynamically:

See “datapath set device policy” on page 389 for more information about the **datapath set device policy** command.

Preferred node path-selection algorithm for the virtualization products

The virtualization products are two-controller disk subsystems. SDD distinguishes the paths to a virtualization product LUN as follows:

1. Paths on the preferred controller
2. Paths on the alternate controller

When SDD selects paths for I/O, preference is always given to a path on the preferred controller. Therefore, in the selection algorithm, an initial attempt is made to select a path on the preferred controller. Only if no path can be used on the preferred controller will a path be selected on the alternate controller. This means that SDD will automatically fail back to the preferred controller any time a path on the preferred controller becomes available during either manual or automatic recovery. Paths on the alternate controller are selected at random. If an error occurs and a path retry is required, retry paths are first selected on the preferred controller. If all retries fail on the preferred controller's paths, paths on the alternate controller will be selected for retry. The following is the path selection algorithm for SDD:

1. With all paths available, I/O is only routed to paths on the preferred controller.

2. If no path on the preferred controller is available, I/O fails over to the alternate controller.
3. When failover to the alternate controller has occurred, if a path on the preferred controller is made available, I/O automatically fails back to the preferred controller.

SDD datapath query adapter command changes for SDD 1.4.0.0 (or later)

For SDD 1.4.0.0 (or later), the output of some of the datapath commands has changed. See Chapter 10, “Using the datapath commands,” on page 371 for details about the datapath commands.

For SDD 1.3.3.11 (or earlier), the output of the **datapath query adapter** command shows all the fibre-channel arrays as different adapters, and you need to determine which hardware paths relate to which adapters. If you need to place an adapter offline, you need to manually issue multiple commands to remove all the associated hardware paths.

For SDD 1.4.0.0 (or later), the output of the **datapath query adapter** command has been simplified.

The following examples show the output resulting from the **datapath query adapter** command for the same configuration for SDD 1.3.3.11 (or earlier) and for SDD 1.4.0.0 (or later).

Example output from **datapath query adapter** command issued in SDD 1.3.3.11 (or earlier):

```
Active Adapters :8
Adapter#      Adapter Name  State   Mode   Select  Error Path Active
0             0/7/0/0.4.18.0.38  NORMAL ACTIVE    0      0    1    1
1             0/4/0/0.4.18.0.38  NORMAL ACTIVE    0      0    1    1
2             0/7/0/0.4.18.0.36  NORMAL ACTIVE    0      0    2    2
3             0/4/0/0.4.18.0.36  NORMAL ACTIVE    0      0    2    2
4             0/7/0/0.4.18.0.34  NORMAL ACTIVE    0      0    2    2
5             0/4/0/0.4.18.0.34  NORMAL ACTIVE    0      0    2    2
6             0/7/0/0.4.18.0.32  NORMAL ACTIVE    0      0    1    1
7             0/4/0/0.4.18.0.32  NORMAL ACTIVE    0      0    1    1
```

Adapter #s 0, 2, 4, 6 belong to the same physical adapter. In order to place this adapter offline, you need to issue **datapath set adapter offline** four times. After the four commands are issued, the output of **datapath query adapter** will be:

```
Active Adapters :8
Adapter#      Adapter Name  State   Mode   Select  Error Path Active
0             0/7/0/0.4.18.0.38  NORMAL OFFLINE   0      0    1    0
1             0/4/0/0.4.18.0.38  NORMAL ACTIVE    0      0    1    0
2             0/7/0/0.4.18.0.36  NORMAL OFFLINE   0      0    2    0
3             0/4/0/0.4.18.0.36  NORMAL ACTIVE    0      0    2    0
4             0/7/0/0.4.18.0.34  NORMAL OFFLINE   0      0    2    0
5             0/4/0/0.4.18.0.34  NORMAL ACTIVE    0      0    2    0
6             0/7/0/0.4.18.0.32  NORMAL OFFLINE   0      0    1    0
7             0/4/0/0.4.18.0.32  NORMAL ACTIVE    0      0    1    0
```

Example output from **datapath query adapter** command issued in SDD 1.4.0.0 (or later):

```
Active Adapters :2
Adapter#      Adapter Name  State   Mode   Select  Error Path Active
0             0/7/0/0      NORMAL ACTIVE    0      0    6    6
1             0/4/0/0      NORMAL ACTIVE    0      0    6    6
```

Adapters 0 and 1 represent two physical adapters. To place one of the adapters offline, you need to issue one single command, for example, **datapath set adapter 0 offline**. After the command is issued, the output of **datapath query adapter** will be:

```
Active Adapters :2
Adapter#      Adapter Name  State   Mode   Select  Error Path Active
0             0/7/0/0     NORMAL OFFLINE 0       0     6     0
1             0/4/0/0     NORMAL ACTIVE  0       0     6     0
```

SDD datapath query device command changes for SDD 1.4.0.0 (or later)

The following change is made in SDD for the **datapath query device** command to accommodate the serial numbers of supported storage devices. The locations of Serial and Policy are swapped because the SAN Volume Controller serial is too long to fit in the first line.

Example output from **datapath query device** command issued in SDD 1.3.3.11 (or earlier):

```
Dev#: 3      Device Name: vpath5  Type: 2105800   Serial: 14123922
Policy:      Optimized
=====
Path#      Adapter H/W Path   Hard Disk      State   Mode   Select  Error
0          0/7/0/0     c19t8d1       OPEN   NORMAL 3869815  0
1          0/7/0/0     c13t8d1       OPEN   NORMAL 3872306  0
2          0/3/0/0     c17t8d1       OPEN   NORMAL 3874461  0
3          0/3/0/0     c11t8d1       OPEN   NORMAL 3872868  0
Dev#: 3      Device Name: vpath5  Type: 2105800   Policy:      Optimized
Serial: 14123922
=====
Path#      Adapter H/W Path   Hard Disk      State   Mode   Select  Error
0          0/7/0/0     c19t8d1       OPEN   NORMAL 3869815  0
1          0/7/0/0     c13t8d1       OPEN   NORMAL 3872306  0
2          0/3/0/0     c17t8d1       OPEN   NORMAL 3874461  0
3          0/3/0/0     c11t8d1       OPEN   NORMAL 3872868  0
```

Example output from **datapath query device** command issued in SDD 1.4.0.0 (or later): (This example shows a SAN Volume Controller and a disk storage system device.)

```
Dev#: 2      Device Name: vpath4  Type: 2145      Policy:      Optimized
Serial: 6005676801850687000000000000000000
=====
Path#      Adapter H/W Path   Hard Disk      State   Mode   Select  Error
0          0/7/0/0     c23t0d0       OPEN   NORMAL 2736767  62
1          0/7/0/0     c9t0d0        OPEN   NORMAL 6         6
2          0/3/0/0     c22t0d0       OPEN   NORMAL 2876312  103
3          0/3/0/0     c8t0d0        OPEN   NORMAL 102      101
Dev#: 3      Device Name: vpath5  Type: 2105800   Policy:      Optimized
Serial: 14123922
=====
Path#      Adapter H/W Path   Hard Disk      State   Mode   Select  Error
0          0/7/0/0     c19t8d1       OPEN   NORMAL 3869815  0
1          0/7/0/0     c13t8d1       OPEN   NORMAL 3872306  0
2          0/3/0/0     c17t8d1       OPEN   NORMAL 3874461  0
3          0/3/0/0     c11t8d1       OPEN   NORMAL 3872868  0
```

Note: vpathname *vpathN* is reserved once it is assigned to a LUN even after the LUN has been removed from the host. The same vpathname, *vpathN*, will be assigned to the same LUN when it is reconnected to the host.

Postinstallation

After SDD is installed, the device driver resides above the HP SCSI disk driver (sdisk) in the protocol stack. In other words, SDD now communicates to the HP-UX device layer. The SDD software installation procedure installs a number of SDD components and updates some system files. Those components and files are listed in Table 17 through Table 19 on page 189.

Table 17. SDD components installed for HP-UX host systems

File	Location	Description
mod.o	/opt/IBMsdd/bin	Object file for the SDD driver module
Executables	/opt/IBMsdd/bin	Configuration and status tools
README.sd	/opt/IBMsdd	README file
sddsrv	/sbin/sddsrv	SDD server daemon
sample_sddsrv.conf	/etc/	Sample SDD server configuration file
sddserver	/sbin/init.d	Script to start or stop the SDD daemon at system up/down time
confserver	/sbin/init.d	Script to load SDD driver and run cfgvpath at system boot time
mvserver	/sbin/init.d	Script to move /sbin/rc1.d/S100localmount to /sbin/rc1.d/S250localmount in order to fix the auto mount problem for SDD vpath device filesystems
datapath.1	/usr/local/man/man1/datapath.1	Manpage for datapath
rmvpath.1	/usr/local/man/man1/rmvpath.1	Manpage for rmvpath
showvpath.1	/usr/local/man/man1/showvpath.1	Manpage for showvpath
gettrace.1	/usr/local/man/man1/gettrace.1	Manpage for gettrace
sddsrv.1	/usr/local/man/man1/sddsrv.1	Manpage for sddsrv
vp2hd.1	/usr/local/man/man1/vp2hd.1	Manpage for vp2hd
hd2vp.1	/usr/local/man/man1/hd2vp.1	Manpage for hd2vp
cfgvpath.1	/usr/local/man/man1/cfgvpath.1	Manpage for cfgvpath
vpcluster.1	/usr/local/man/man1/vpcluster.1	Manpage for vpcluster
sddgetdata.1	/usr/local/man/man1/sddgetdata.1	Manpage for sddgetdata

Table 18. System files updated for HP-UX host systems

File	Location	Description
vpath	/usr/conf/master.d	Master configuration file
vpath	/stand/system.d	System configuration file

Table 19. SDD commands and their descriptions for HP-UX host systems

Command	Description
cfgvpath	Configures new SDD vpath devices when there are no existing vpath devices. Do not use the legacy parameter "-c".
cfgvpath [-r] (Dynamic reconfiguration)	The SDD vpath device configuration is updated without system reboot. If initially there is no SDD vpath device configured, cfgvpath -r will fail with message "failed to get information from kernel, don't run dynamic configuration, do cfgvpath instead." In this case, issue cfgvpath without any option.
showvpath	Lists the configuration mapping between the SDD devices and underlying disks.
datapath	The SDD driver console command tool.
hd2vp	Converts a volume group from sdisks into the SDD vpath devices.
vp2hd	Converts a volume group from the SDD vpath devices into sdisks.
vpcluster	Imports or exports MC Service Guard volume groups.
rmvpath [-all, -vpathname]	Removes the SDD vpath devices from the configuration.
gettrace	Debug tool that gets trace information when problem occurs.
sddgetdata	The SDD data collection tool for problem analysis.
man	Manpage for SDD commands, for example, man datapath . Supported SDD commands are datapath , gettrace , hd2vp , querysn , rmvpath , sddsrv , sddgetdatashowvpath , vp2hd , vpcluster , and sddgetdata .

If you are not using a DBMS or an application package that communicates directly to the sdisk interface, the installation procedure is nearly complete. However, you still need to customize HP-UX so that standard UNIX[®] applications can use the SDD. Go to "Standard UNIX applications" on page 196 for instructions. If you have a DBMS or an application package installed that communicates directly to the sdisk interface, such as Oracle, go to "Using applications with SDD" on page 196 and read the information specific to the application that you are using.

During the installation process, the following files were copied from the IBMsdd_depot to the system:

Kernel-related files

- /opt/IBMsdd/bin/mod.o
- /stand/system.d/vpath
- /usr/conf/master.d/vpath

SDD driver-related files

- /opt/IBMsdd

- /opt/IBMsdd/bin
- /opt/IBMsdd/README.sd
- /opt/IBMsdd/bin/cfgvpath
- /opt/IBMsdd/bin/datapath
- /opt/IBMsdd/bin/showvpath
- /opt/IBMsdd/bin/master
- /opt/IBMsdd/bin/system
- /opt/IBMsdd/bin/mod.o
- /opt/IBMsdd/bin/rmvpath
- /opt/IBMsdd/bin/get_root_disks
- /opt/IBMsdd/bin/gettrace
- /opt/IBMsdd/bin/sddgetdata
- /opt/IBMsdd/bin/hd2vp
- /opt/IBMsdd/bin/vp2hd
- /opt/IBMsdd/bin/vpcluster
- /sbin/cfgvpath
- /sbin/datapath
- /sbin/get_root_disks
- /sbin/rmvpath
- /sbin/showvpath
- /sbin/hd2vp
- /sbin/vp2hd
- /sbin/vpcluster
- /sbin/sddgetdata
- /sbin/sddsrv
- /etc/sample_sddsrv.conf

During installation, /opt/IBMsdd/bin/cfgvpath program is initiated to create SDD vpath devices in the /dev/dsk and /dev/rdisk directories for all IBM disks that are available on the system. After installation is done, all SDD vpath devices are configured and the driver is loaded. The system will not reboot.

Note: SDD devices are found in /dev/rdisk and /dev/dsk. The device is named according to the SDD number. A device with a number of 0 would be /dev/rdisk/vpath1.

Uninstalling SDD

The following procedure explains how to remove SDD. You must uninstall the current level of SDD before upgrading to a newer level.

Complete the following procedure to uninstall SDD:

1. Stop applications.
2. If you are using SDD with a database, such as Oracle, edit the appropriate database configuration files (database partition) to remove all the SDD devices.
3. Before running the sam program, run script **vp2hd** to convert volume groups from SDD vpath devices to sdisks.
4. Run the **sam** program.

- > **sam**
- 5. Click **Software Management**.
- 6. Click **Remove Software**.
- 7. Click **Remove Local Host Software**.
- 8. Click the **IBMsdd_tag** selection.
 - a. From the Bar menu, click **Actions** → **Mark for Remove**.
 - b. From the Bar menu, click **Actions** → **Remove (analysis)**. A Remove Analysis window opens and shows the status of Ready.
 - c. Click **OK** to proceed. A Confirmation window opens and indicates that the uninstallation will begin.
 - d. Click **Yes**. The analysis phase starts.
 - e. After the analysis phase has finished, another Confirmation window opens indicating that the system will be restarted after the uninstallation is complete. Click **Yes** and press **Enter**. The uninstallation of IBMsdd begins.
 - f. An Uninstall window opens showing the progress of the IBMsdd software uninstallation. This is what the panel looks like:

```

Target      : XXXXX
Status      : Executing unconfigure
Percent Complete : 17%
Kbytes Removed : 340 of 2000
Time Left (minutes) : 5
Removing Software : IBMsdd_tag,.....
  
```

The **Done** option is not available when the uninstallation process is in progress. It becomes available after the uninstallation process completes.

- 9. Click **Done**.

When SDD has been successfully uninstalled, the first part of the procedure for upgrading the SDD is complete. To complete an upgrade, you need to reinstall SDD. See the installation procedure in “Installing SDD” on page 179.

Note: The MC Service Guard cluster must be halted prior to uninstall. The primary node and the adoptive node or nodes must operate in a single-host environment. The shared volume groups in the adoptive nodes must be exported so that the volumes are not shared; the volume groups can be active in the primary node only.

SDD server daemon

The SDD server (also referred to as `sddsrv`) is an integrated component of SDD 1.3.1.5 (or later). This component consists of a UNIX application daemon that is installed in addition to the SDD device driver. See Chapter 9, “Using the SDDDSM or the SDDPCM daemon,” on page 365 for more information about `sddsrv`.

Note: The SDD server is started from `/etc/inittab`. Because the environment variable `$TZ` is not exported at init time, you should modify `/etc/default/tz` to reflect the corresponding time zone so that `sddsrv.log` reflects the correct local time. See the man pages for the `ctime` command for more information.

Verifying if the SDD server has started

After you have installed SDD, verify that the SDD server (`sddsrv`) has automatically started by entering `ps -ef | grep sddsrv`.

If the SDD server (sddsrv) has automatically started, the output will display the process number on which sddsrv has started.

If sddsrv has *not* started, you should uninstall SDD and then reinstall SDD. See “Installing SDD” on page 179 for more information.

Starting the SDD server manually

If the SDD server does not start automatically after you complete the SDD installation, or if you want to start it manually after stopping sddsrv, use the following process to start sddsrv:

1. Edit /etc/inittab and verify the sddsrv entry.

For example:

```
srv:23456:respawn:/sbin/sddsrv >/dev/null 2>&1
```

2. Save the file /etc/inittab.
3. Issue **init q**.

Go to “Verifying if the SDD server has started” on page 191 for the steps to see if you successfully started the SDD server.

Changing to a different port number for the SDD server

To change to a different port number for the SDD server, see the instructions in “Changing the sddsrv or pcmsrv TCP/IP port number” on page 369.

Stopping the SDD server

To stop the SDD server, complete the following steps:

1. Edit /etc/inittab and comment out the SDD server entry:

```
#srv:23456:respawn:/sbin/sddsrv >/dev/null 2>&1 2
```

2. Save the file.
3. Issue **init q**.
4. Check if sddsrv is running by issuing **ps -ef | grep sddsrv**. If sddsrv is still running, issue **kill -9 pid** of sddsrv.

How to import and export volume groups

Use the HP CLI to manage SDD devices.

You can import volume groups that are created over SDD vpath devices using the **vgimport** command. The **vgimport** command is useful in conjunction with the **vgexport** command.

Before you can import the specified volume groups, you must complete the following tasks:

1. Export or move volume groups from one node to another node within a high availability cluster by using the **vgexport** command. See “Exporting volume groups” on page 193.
2. FTP the map file to the other node within a high-availability cluster. See “Moving the map file” on page 193.
3. Create the volume group device directory. See “Creating the volume group device directory” on page 193.

4. Create the group special file. See “Creating the group special file” on page 194. For more information about the **vgimport** command, see “Importing volume groups” on page 194.

Exporting volume groups

The **vgexport** command recognizes the following options and arguments:

- p** The **-p** option previews the actions to be taken but does not update the `/etc/lvmtab` file or remove the devices file.
- v** The **-v** option prints verbose messages including the names of the physical volumes associated with this volume group.
- s** **-s** is the sharable option (Series 800 only). When you specify the **-s** option, you must also specify the **-p**, **-v**, and **-m** options. A mapfile is created that can be used to create volume group entries (with the **vgimport** command) on other systems in the high availability cluster.

-m mapfile

By default, a file named `mapfile` is created in your current directory. The mapfile contains a description of the volume group and its associated logical volumes. Use the **-m** option to specify a different name for the mapfile. The mapfile serves as input to **vgimport**; When the mapfile is used with the **-s** option, the volume group specified in the mapfile can be shared with the other systems in the high availability cluster.

vg_name

The `vg_name` is the path name of the volume group.

vgexport command example:

To export the specified volume group on node 1, enter:

```
vgexport -p -v -s -m /tmp/vgpath1.map vgpath1
```

where `/tmp/vgpath1.map` represents your mapfile, and `vgpath1` represents the path name of volume group that you want to export.

Moving the map file

You must also FTP the map file to the other node.

For example, to FTP the `vgpath1.map` map file to node 2, enter:

```
rcp /tmp/vgpath1.map node2:/tmp/vgpath1.map
```

Creating the volume group device directory

You must also create the volume group device directory.

For example, to create the volume group device directory `/dev/vgpath1` on node 2, enter:

```
mkdir /dev/vgpath1
```

Creating the group special file

You must also create the group special file on node 2.

For example, to create the *group c 64* file, enter:

```
mknod /dev/vgpath1/group c 64 n
```

where *n* is the same as that was given when */dev/vgpath1/group* was created on node 1.

Importing volume groups

The **vgimport** command recognizes the following options and arguments:

- p** The **-p** option previews the actions to be taken but does not update the */etc/lvmtab* file or remove the devices file.
- v** The **-v** option prints verbose messages including the names of the logical volumes.
- s** **-s** is the sharable option (disk storage system Series 800 only). When you specify the **-s** option, you must specify the **-p**, **-v**, and **-m** options. The specified mapfile is the same mapfile specified by using the **vgexport** command also using the **-p**, **-m**, and **-s** options. The mapfile is used to create the volume groups on the importing systems.

-m mapfile

By default, a file named *mapfile* is created in your current directory. The mapfile contains a description of the volume group and its associated logical volumes. Use the **-m** option to specify a different name for the mapfile. The mapfile serves as input to **vgimport**; When the mapfile is used with the **-s** option, the volume group specified in the mapfile can be shared among the exporting system and the importing system.

vg_name

The *vg_name* is the path name of the volume group.

vgimport command example:

To import the specified volume group on node 2, enter:

```
vgimport -p -v -s -m /tmp/vgpath1.map vgpath1
```

where */tmp/vgpath1.map* represents your mapfile, and *vgpath1* represents the path name of the volume group that you want to import.

Note: The **vgimport** command only imports the scsi pmlink devices. Issue the **hd2vp** command after issuing the **vgimport** command.

Exporting or importing MC Service Guard volume groups

The **vpcluster** script makes it easier to export or import MC Service Guard volume groups managed by SDD. The normal Service Guard configuration and SDD device configuration process remain the same. The sharable volume groups used by the HA clusters must be marked

For the primary node operation, **vpcluster** extracts the node and volume group information from the keywords *NODE_NAME* and *VOLUME_GROUP* as specified in the cluster configuration file. It ensures that the devices of these volumes groups

are SDD-managed vpath devices, performs **vgexport** and creates **vpcluster** control files for the adoptive nodes to import. The input file to **vpcluster** does not have to be the same cluster configuration file for the SG cluster creation. It can be any ASCII file with the line entries that contain **NODE_NAME** and **VOLUME_GROUP** keywords without # as a comment. Optionally, the **vpcluster** control file can be copied to each adoptive node with the **rcp** command.

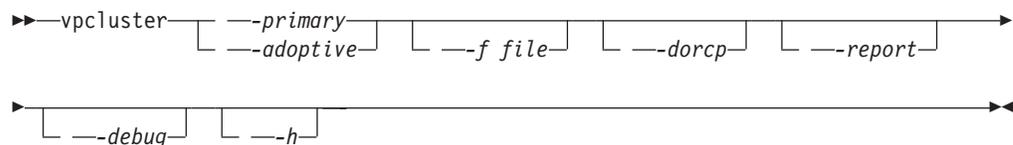
For the adoptive node operation, **vpcluster** uses the control file created by the primary node operation. Prior to volume group import, it validates that the adoptive node is included in the cluster configuration, ensures the importing volume groups are not active volume groups in the adoptive node, creates volume group nodes `/dev/vgXXXX` using the **mknod** command, and ensures that the same device serial and LUN-id are configured by SDD.

Note:

1. The device names, either `vpath#` or `C#T#D#`, might be different between the primary and adoptive nodes. However, the **vpcluster** process attempts to keep the volume group minor number consistent between the primary and adoptive nodes. In case the same minor number is in use, the next sequential number is assigned. Because the HP **vgimport** process only imports those device names with the standard name `C#T#D#`, **hd2vp** is invoked to convert `sdisk` devices to SDD devices after a successful **vgimport**.
2. Use the **cmquerycl** HP command to create the cluster configuration file. This command recognizes only the `pmlink scsi` devices on both nodes. If you are using the **cmquerycl** command to create the cluster configuration file, you should first issue **vp2hd** to convert `vpath` devices to `pmlink` devices.

In addition, the `report` option for adoptive node validates all volume groups exported by the primary node. A mismatch of volume group minor number or `vpath` device name is allowed. Other mismatches will be reported.

Syntax:



where,

-primary

Specifies primary node operation. You must specify **-primary** or **-adoptive**.

-adoptive

Specifies adoptive node operation. You must specify **-primary** or **-adoptive**.

-f file

For the primary node, specify the cluster configuration file, default is `/etc/cmcluster/cmclconf.ascii`.

For the adoptive node, specify the **vpcluster** control file created the primary node, default is `/tmp/vpcluster/vpcluster.primary.tar`

-dorcp

Specifies the **vpcluster** control tar file to be RCPed to the adoptive nodes. The default is no.

-report

Validates that the exported volume groups from the primary node are imported to the adoptive node and creates a report. This option is valid in the adoptive node.

-debug

Specifies that a debugging statement is to be printed during **vpcluster** run time.

-h Specifies that detailed help info about the **vpcluster** function is to be displayed.

There are more than one way to configure the SG cluster locking: quorum server, or lock disk. In case the lock disk is chosen, the SDD vpath device should not be used since it is not recognized by the `FIRST_CLUSTER_LOCK_PV` parameter. Furthermore, it is recommended that SDD vpath devices and `sdisk` pmlink should not be mixed in the same volume group. The lock device should be excluded from the SDD configuration. See the information about the `/etc/vpathmanualexcl.cfg` text file on page 184.

Using applications with SDD

If your system already has a software application or a DBMS installed that communicates directly with the HP-UX disk device drivers, you need to insert the new SDD device layer between the software application and the HP-UX disk device layer. You also need to customize the software application to have it communicate with the SDD devices instead of the HP-UX devices.

In addition, many software applications and DBMSs need to control certain device attributes such as ownership and permissions. Therefore, you must ensure that the new SDD devices that these software applications or DBMSs access in the future have the same attributes as the HP-UX `sdisk` devices that they replace. You need to customize the application or DBMS to accomplish this.

This section contains the procedures for customizing the following software applications and DBMSs for use with SDD:

- Standard UNIX applications
- Network File System (NFS) file server

Standard UNIX applications

Note: The SDD path state change depends on the “positive” feedback from an I/O request. The positive feedback is an I/O response, either successful or failed condition, from an I/O request. Due to the standard and designed behavior of the SCSI driver on HP-UX 11i V2 and prior, the following conditions occur when access raw `sdisk` through SDD:

1. I/O errors are retried indefinitely because pre-HP-UX 11i v3 (B.11.31) has no way to distinguish between a non-volume-manager file system I/O and a raw access.
2. I/O errors returned on certain file system I/Os can result in panic.

SDD can still load balance with the raw device access but failover/failback is not possible. In most cases, an I/O error is likely to hang due to indefinite retries at the SCSI driver layer.

If you have not already done so, install SDD using the procedure in “Installing SDD” on page 179. When this is done, SDD resides above the HP-UX SCSI disk

driver (sdisk) in the protocol stack. In other words, SDD now communicates to the HP-UX device layer. To use standard UNIX applications with SDD, you must make some changes to your logical volumes. You must convert your existing logical volumes or create new ones.

Standard UNIX applications such as `newfs`, `fsck`, `mkfs`, and `mount`, which normally take a disk device or raw disk device as a parameter, also accept the SDD device as a parameter. Similarly, entries in files such as `vfstab` and `dfstab` (in the format of `cntndnsn`) can be replaced by entries for the corresponding SDD `vpathNs` devices. Make sure that the devices that you want to replace are replaced with the corresponding SDD device. Issue the `showvpath` command to list all SDD `vpath` devices and their underlying disks.

To use the SDD driver for an existing logical volume, you must run the `hd2vp` conversion script (see “SDD utility programs” on page 86).

Attention: Do not use the SDD for critical file systems needed at startup, such as `/(root)`, `/stand`, `/usr`, `/tmp` or `/var`. Doing so may render your system unusable if SDD is ever uninstalled (for example, as part of an upgrade).

Creating new logical volumes

Use the following process to create a new logical volume to use SDD:

Note: You must have superuser privileges to complete these subtasks.

1. Determine the major number of the logical volume device.

Enter the following command to determine the major number:

```
# lsdev | grep lv
```

A message similar to the following is displayed:

```
64          64          1v          1vm
```

The first number in the message is the major number of the character device, which is the number that you want to use.

2. Create a device node for the logical volume device.

Note: If you do not have any other logical volume devices, you can use a minor number of `0x010000`. In this example, assume that you have no other logical volume devices. A message similar to the following is displayed:

```
# mknod group c 64 0x010000
```

Create a physical volume by performing the procedure in step 3 on page 198.

- a. Create a subdirectory in the `/dev` directory for the volume group.

Enter the following command to create a subdirectory in the `/dev` directory for the volume group:

```
# mkdir /dev/vgIBM
```

In this example, `vgIBM` is the name of the directory.

- b. Change to the `/dev` directory.

Enter the following command to change to the `/dev` directory:

```
# cd /dev/vgIBM
```

- c. Create a device node for the logical volume device.
Enter the following command to re-create the physical volume:

```
# pvcreate /dev/rdisk/vpath1
```

A message similar to the following is displayed:

```
Physical volume "/dev/rdisk/vpath1" has been successfully created.
```

In this example, the SDD vpath device associated with the underlying disk is vpath1. Verify the underlying disk by entering the following **showvpath** command:

```
# /opt/IBMsdd/bin/showvpath
```

A message similar to the following is displayed:

```
vpath1:  
/dev/dsk/c3t4d0
```

3. Create a physical volume.
Enter the following command to create a physical volume:

```
# pvcreate /dev/rdisk/vpath1
```

4. Create a volume group.
Enter the following command to create a volume group:

```
# vgcreate /dev/vgIBM/dev/dsk/vpath1
```

5. Create a logical volume.
Enter the following command to create logical volume *lvol1*:

```
# lvcreate -L 100 -n lvol1 vgIBM
```

The `-L 100` portion of the command makes a 100-MB volume group; you can make it larger if you want to. Now you are ready to create a file system on the volume group.

6. Create a file system on the volume group.
Use the following process to create a file system on the volume group:
 - a. If you are using an HFS file system, enter the following command to create a file system on the volume group:

```
# newfs -F HFS /dev/vgIBM/rlvol1
```

- b. If you are using a VXFS file system, enter the following command to create a file system on the volume group:

```
# newfs -F VXFS /dev/vgIBM/rlvol1
```

- c. Mount the logical volume.
This process assumes that you have a mount point called `/mnt`.
7. Mount the logical volume.
Enter the following command to mount the logical volume *lvol1*:

```
# mount /dev/vgIBM/lvol1 /mnt
```

Attention: In some cases, it may be necessary to use standard HP-UX recovery procedures to fix a volume group that has become damaged or corrupted. For information about using recovery procedures, such as **vgscan**, **vgextend**, **vpchange**, or **vgreduce**, search for *Manuals for HP-UX 11i Volume Management (LVM/VxVM) Software* in the following website:

h20565.www2.hp.com/portal/site/hpsc/public/

Removing logical volumes

Use the following procedure to remove logical volumes:

1. Remove the existing logical volume.

Before the logical volume is removed, it must be unmounted. For example, enter the following command to unmount logical volume *lv01*:

```
# umount /dev/vgIBM/lv01
```

Next, remove the logical volume.

For example, enter the following command to remove logical volume *lv01*:

```
# lvremove /dev/vgIBM/lv01
```

A message similar to the following is displayed:

```
The logical volume "/dev/vgIBM/lv01" is not empty;
do you really want to delete the logical volume (y/n)
```

Enter **y** and press **Enter**. A message similar to the following is displayed:

```
Logical volume "/dev/vgIBM/lv01" has been successfully removed.
Volume Group configuration for /dev/vgIBM has been saved in
/etc/lvmconf/vgIBM.conf
```

When prompted to delete the logical volume, enter **y**.

2. Remove the existing volume group.

Enter the following command to remove the volume group *vgIBM*:

```
# vgrename /dev/vgIBM
```

A message similar to the following is displayed:

```
Volume group "/dev/vgIBM" has been successfully removed.
```

Now, you can re-create the logical volume.

Re-creating the existing logical volumes

Use the following process to convert an existing logical volume to use SDD:

Note: You must have superuser privileges to complete these subtasks.

As an example, suppose you have a logical volume called *lv01* under a volume group *vgIBM*, which is currently using the disk directly, (for example, through path */dev* path */dev/dsk/c3t4d0*). You want to convert logical volume *lv01* to use SDD.

1. Determine the size of the logical volume.

Enter the following command to determine the size of the logical volume:

```
# lvsdisplay /dev/vgIBM/lvol1 | grep "LV Size"
```

A message similar to the following is displayed:

```
LV Size (Mbytes) 100
```

In this case, the logical volume size is 100 MB.

2. Re-create the physical volume.

Enter the following command to re-create the physical volume:

```
# pvcreate /dev/rdisk/vpath1
```

A message similar to the following is displayed:

```
Physical volume "/dev/rdisk/vpath1" has been successfully created.
```

In this example, the SDD vpath device associated with the underlying disk is vpath1. Verify the underlying disk by entering the following command:

```
# /opt/IBMsdd/bin/showvpath
```

A message similar to the following is displayed:

```
vpath1:  
/dev/dsk/c3t4d0
```

3. Re-create the volume group.

Enter the following command to re-create the volume group:

```
# vgcreate /dev/vgibm /dev/dsk/vpath1
```

A message similar to the following is displayed:

```
Increased the number of physical extents per physical volume to 2187.  
Volume group "/dev/vgibm" has been successfully created.  
Volume Group configuration for /dev/vgibm has been saved in  
/etc/lvmconf/vgibm.conf
```

4. Re-create the logical volume.

Re-creating the logical volume consists of a number of smaller steps:

- a. Re-creating the physical volume
- b. Re-creating the volume group
- c. Re-creating the logical volume

Enter the following command to re-create the logical volume:

```
# lvcreate -L 100 -n lvol1 vgibm
```

A message similar to the following is displayed:

```
Logical volume "/dev/vgibm/lvol1" has been successfully created with  
character device "/dev/vgibm/r1vol1".  
Logical volume "/dev/vgibm/lvol1" has been successfully extended.  
Volume Group configuration for /dev/vgibm has been saved in  
/etc/lvmconf/vgibm.conf
```

The -L 100 parameter comes from the size of the original logical volume, which is determined by using the **lvdisplay** command. In this example, the original logical volume was 100 MB in size.

Attention: The re-created logical volume should be the same size as the original volume; otherwise, the re-created volume cannot store the data that was on the original.

5. Setting the proper timeout value for the logical volume manager.

The timeout values for the Logical Volume Manager must be set correctly for SDD to operate properly. This is particularly true if the concurrent firmware download has taken place. There are two timeout values: one for logical volume (LV) and one for physical volume (PV).

The LV timeout value is determined by the application. If the application has no specific timeout requirement, use the HP default value, which is 0 (forever).

The PV timeout value is recommended by the storage vendor. The HP default PV timeout value is 30 seconds. Generally, this is sufficient during normal operations.

However, during the concurrent firmware download, you must set the PV timeout value to a minimum of 90 seconds. You can set the timeout value to 90 seconds for normal operations as well. In addition, if you do not use the default LV timeout value, ensure that the LV timeout value is no less than the PV timeout value multiplied by the number of paths. For example, when the default is not used, if a vpath device has four underlying paths and the PV timeout value is 90, the LV timeout value must be at least 360.

To display the timeout value, use the **lvdisplay** or **pvdisplay** command.

To change the PV timeout value, use the **pvchange** command after **pvcreate**, and to change the LV timeout value, use the **lvchange** command after **lvcreate**.

For example:

- To change the timeout value of all underlying paths of vpathX to 90 seconds, enter **pvchange -t 90 /dev/dsk/vpathX**
- To change the timeout value for logical volume /dev/vgibm/lvolY to 360 seconds, enter **lvchange -t 360 /dev/vgibm/lvolY**

Attention: In some cases, it may be necessary to use standard HP-UX recovery procedures to fix a volume group that has become damaged or corrupted. For information about using recovery procedures, such as **vgscan**, **vgextend**, **pvchange**, or **vgreduce**, search for *Manuals for HP-UX 11i Volume Management (LVM/VxVM) Software* in the following website:

h20565.www2.hp.com/portal/site/hpsc/public/

Installing SDD on a NFS file server

The procedures in this section show how to install SDD for use with an exported file system (NFS file server).

Setting up NFS for the first time

Complete the following steps if you are installing exported file systems on SDD devices for the first time:

1. If you have not already done so, install SDD using the procedure in “Installing SDD” on page 179.
2. Determine which SDD (vpathN) volumes that you will use as file system devices.

3. Create file systems on the selected SDD devices using the appropriate utilities for the type of file system that you will use. If you are using the standard HP-UX UFS file system, enter the following command:

```
# newfs /dev/rdisk/vpathN
```

In this example, *N* is the SDD device instance of the selected volume. Create mount points for the new file systems.

4. Install the file systems into the directory `/etc/fstab`. In the **mount at boot** field, click **yes**.
5. Install the file system mount points into the `/etc/exports` directory for export.
6. Restart the system.

Installing SDD on a system that already has the NFS file server

Complete the following steps if you have the NFS file server already configured to:

- Export file systems that reside on a multiport subsystem, and
 - Use SDD partitions instead of `sdisk` partitions to access them
1. List the mount points for all currently exported file systems by looking in the `/etc/exports` directory.
 2. Match the mount points found in step 1 with `sdisk` device link names (files named `/dev/(r)dsk/cntndn`) by looking in the `/etc/fstab` directory.
 3. Match the `sdisk` device link names found in step 2 with SDD device link names (files named `/dev/(r)dsk/vpathN`) by issuing the **showvpath** command.
 4. Make a backup copy of the current `/etc/fstab` file.
 5. Edit the `/etc/fstab` file, replacing each instance of an `sdisk` device link named `/dev/(r)dsk/cntndn` with the corresponding SDD device link.
 6. Restart the system.
 7. Verify that each exported file system:
 - a. Passes the start time **fsck pass**
 - b. Mounts properly
 - c. Is exported and available to NFS clients

If there is a problem with any exported file system after completing step 7, restore the original `/etc/fstab` file and restart to restore NFS service. Then review your steps and try again.

Chapter 5. Using SDD on a Linux host system

This chapter provides step-by-step procedures on how to install, configure, use, and remove SDD on supported Linux host systems that are attached to supported storage devices. For updated and additional information that is not included in this chapter, see the Readme file on the CD-ROM or visit the SDD website:

www-01.ibm.com/support/docview.wss?uid=ssg1S4000107

Verifying hardware and software requirements

You must install the following hardware and software components to ensure that SDD installs and operates successfully.

Hardware

The following hardware components are needed:

- Supported storage devices
- One or more pairs of fibre-channel host adapters
To use SDD's input/output (I/O) load balancing features and failover features, you need a minimum of two paths to your storage devices.
For more information about the fibre-channel adapters that you can use on your Linux host system, see the *Host Systems Attachment Guide* for your product.
- Subsystem LUNs that have been created and configured for multiport access. Subsystem LUNs are known as SDD vpath devices in Linux SDD. Each SDD vpath device can have up to 32 paths (SCSI disk instances).
- A fibre optic cable to connect each fibre-channel adapter to a supported storage device port, or to switch ports subsequently zoned to supported storage device ports.

See the *IBM TotalStorage Enterprise Storage Server: Interoperability Guide* for more information regarding hardware, software, and driver support.

Software

A general list of supported Linux distributions and major release levels is shown below. For the most up-to-date information regarding support for specific architectures and kernels, see the Readme file for the latest SDD release on the CD-ROM or visit the SDD website:

www-01.ibm.com/support/docview.wss?uid=ssg1S4000107

- Novell SUSE
 - SUSE Linux Enterprise Server (SLES) 8 / UnitedLinux 1.0
 - SLES 9
- Red Hat
 - RHEL 3 AS
 - RHEL 4 AS
- Asianux
 - Red Flag Advanced Server 4.1
 - Red Flag DC Server 4.1

Unsupported environments

SDD does not support environments containing the following functions:

- DS8000 and DS6000 do not support SCSI connectivity. ESS Model 800 does support SCSI connectivity.
- The EXT3 file system on an SDD vpath device is **only** supported on distributions running the 2.4.21 or newer kernel.
- Single-path mode during concurrent download of licensed machine code nor during any disk storage system concurrent maintenance that impacts the path attachment, such as a disk storage system host-bay-adapter replacement, or host zoning reconfiguration that affects the host or storage ports in use.

Preparing for SDD installation

Before you install SDD, you must configure the supported storage device for your host system and attach required fibre-channel adapters.

Configuring disk storage systems

Before you install SDD, configure your disk storage system for multiport access for each LUN. SDD requires a minimum of two paths to your storage devices that share the same LUN to use the load-balancing and path-failover-protection features. **With a single-path, failover protection is not provided.**

A host system with a single fibre-channel adapter connected through a switch to multiple disk storage system ports is considered a multipath fibre-channel connection.

See the *Planning Guide* and *Host Systems Attachment Guide* for your product to find more information about how to configure the disk storage system. These publications are listed in “Related information” on page xiv.

See the *IBM TotalStorage Enterprise Storage Server Host Systems Attachment Guide* for information on working around Linux LUN limitations.

Configuring virtualization products

Before you install SDD, configure your virtualization product for multiport access for each LUN. SDD requires a minimum of two paths to your storage devices that share the same LUN to use the load-balancing and path-failover-protection features. **With a single-path, failover protection is not provided.**

A host system with a single fibre-channel adapter connected through a switch to multiple disk storage system ports is considered a multipath fibre-channel connection.

For information about configuring your SAN Volume Controller, see the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

See the *IBM System Storage SAN Volume Controller Host Systems Attachment Guide* for information on working around Linux LUN limitations.

Configuring fibre-channel adapters on disk storage systems

You must configure the fibre-channel adapters and the adapter drivers that are attached to your Linux host system before you install SDD. Follow the adapter-specific configuration instructions to configure the adapters.

See the *IBM TotalStorage Enterprise Storage Server Host Systems Attachment Guide* for more information about how to install and configure fibre-channel adapters for your Linux host system and for information about working around Linux LUN limitations.

Configuring fibre-channel adapters on virtualization products

You must configure the fibre-channel adapters and the adapter drivers that are attached to your Linux host system before you install SDD. Follow the adapter-specific configuration instructions to configure the adapters.

For information about configuring your SAN Volume Controller, see the *IBM System Storage SAN Volume Controller Planning Guide*, and the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*. See the *IBM System Storage SAN Volume Controller Host Systems Attachment Guide* for more information about how to install and configure fibre-channel adapters for your Linux host system and for information about working around Linux LUN limitations.

Disabling automatic Linux system updates

Many Linux distributions give you the ability to configure your systems for automatic system updates. Red Hat provides this ability in the form of a program called **up2date**, while Novell SUSE provides the YaST Online Update utility. These features periodically query for updates available for each host and can be configured to automatically install any new updates that they find. Often, the automatic update process also upgrades the system to the latest kernel level. If this is the case, hosts running SDD should consider turning off the automatic update of kernel levels. Some drivers supplied by IBM, such as 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 in order to function optimally. By allowing automatic updates of the kernel, you risk impacting your host systems unexpectedly.

Installing SDD

Before you install SDD, make sure that you have root access to your Linux host system and that all the required hardware and software is ready.

Installing SDD from CD-ROM

Note: A new feature of SDD 1.6.1.x and later packages is the possibility of RPM relocation. RPM relocation allows you to manually relocate the files that SDD installs from the default subdirectory (/opt) to another subdirectory. You can specify the new subdirectory under the root filesystem (/) by using the --prefix command. If you want to use the relocation command, see step 7 below. Relocation cannot be done directly onto the root filesystem (/), but must be a subdirectory or other filesystem under root.

Complete the following steps to install SDD on your Linux host system:

1. Log on to your host system as the root user.
2. Insert the SDD installation CD-ROM into your CD-ROM drive.
3. Enter **mount /dev/cdrom** to mount the CD-ROM drive.
4. Enter the following to access your CD-ROM contents :
 - For Red Hat or Asianux: enter **cd /mnt/cdrom**

- For SUSE: enter `cd /media/cdrom`
5. If you are running Red Hat, enter `cd redhat`. If you are running SUSE, enter `cd suse`, and then enter `ls` to display the name of the package. If you are running Miracle Linux, Red Flag, or Asianux, enter `cd asianux`.
 6. Enter `rpm -qpl IBMsdd-N.N.N.N-x.arch.distro.rpm` to view all the files in the package, where:
 - *N.N.N.N-x* represents the current version release modification level number; for example, *N.N.N.N-x* = 1.6.0.1-1.
 - *arch* is the architecture (i686, ppc64, ia64)
 - *distro* is one of the following:
 - rhel3
 - rhel4
 - ul1
 - sles8
 - sles9
 - asianux
 7. Enter the following command to install SDD: `rpm -ivh [--prefix=newpath] IBMsdd-N.N.N.N-x.arch.distro.rpm` where, *newpath* is the new directory under which you want to place SDD files (the default is /opt). Note that you cannot specify `--prefix=.` The prefix flag is optional.

A message similar to the following is displayed:

```
Preparing for installation ...
IBMsdd-N.N.N.N-1
```

Installing SDD from downloaded code

Use this procedure to install SDD from downloaded code:

1. Download code from the website that is appropriate for your system.
2. Change directory to the directory that contains the downloaded .rpm file.
3. Issue the following instruction:

```
rpm -ivh [--prefix=newpath] IBMsdd-N.N.N.N-x.arch.distro.rpm
```

where *newpath* is the new directory where you want to place SDD files (the default directory is /opt). You cannot specify `--prefix=.` The prefix flag is optional.

Upgrading SDD

Complete the following steps to upgrade SDD on your Linux host system:

1. Log on to your host system as the root user.
2. Insert the SDD installation CD into your CD-ROM drive.
3. Enter `mount /dev/cdrom` to mount the CD-ROM drive.
4. Enter the following to access your CD-ROM contents :
 - For Red Hat or Asianux: enter `cd /mnt/cdrom`
 - For SUSE: enter `cd /media/cdrom`

5. If you are running Red Hat, enter **cd redhat**; if you are running SUSE, enter **cd suse**, and then enter **ls** to display the name of the package.
6. Enter **rpm -qpl IBMsdd-N.N.N.N-x.arch.distro.rpm** to view all the files in the package.
7. Enter **rpm -U IBMsdd-N.N.N.N-x.arch.distro.rpm [--prefix=newpath]** to upgrade SDD. The **--prefix** option should be used if it was used during the RPM installation of SDD.

A message similar to the following is displayed:

```
Preparing for installation ...
IBMsdd-N.N.N.N-1
```

Note: The RPM upgrade command (**rpm -U**) will not work if you want to upgrade from a pre-SDD 1.6.1.x package to an SDD 1.6.1.x or later package. Instead:

1. Uninstall the SDD package using the RPM erase command (**rpm -e IBMsdd**)
2. Install the new SDD 1.6.1.x or later package using **rpm -i**.
3. If you modified your `/etc/vpath.conf`, the **rpm -e** command saved a copy in `/etc/vpath.conf.rpmsave`. To preserve your `/etc/vpath.conf` modifications, you must also copy the `/etc/vpath.conf.rpmsave` to `/etc/vpath.conf`.

Verifying the SDD installation

The SDD installation installs the device driver and its utilities in the `/opt/IBMsdd` directory. Table 20 lists the SDD driver and its major component files.

Table 20. SDD components for a Linux host system

File name	Location ¹	Description
sdd-mod.o-xxx (for Linux 2.4 and earlier kernels)	/opt/IBMsdd	SDD device driver file (where XXX stands for the kernel level of your host system).
sdd-mod.ko-xxx (For Linux 2.6 kernels only)	/opt/IBMsdd	SDD device driver file (where XXX stands for the kernel level of your host system).
vpath.conf	/etc	SDD configuration file
sddsrv.conf	/etc	sddsrv configuration file
executables	/opt/IBMsdd/bin	SDD configuration and status tools
	/usr/sbin	Symbolic links to the SDD utilities
sdd.rcscript	/etc/init.d/sdd	Symbolic link for the SDD system startup option
	/usr/sbin/sdd	Symbolic link for the SDD manual start or restart option

¹ In this table, the `/opt` directory is the default directory. The root prefix might be different, depending on the installation.

You can issue the **rpm -qi IBMsdd** command to receive information on the particular package, or **rpm -ql IBMsdd** command to list the specific SDD files that were successfully installed on your Linux host system. If the installation was successful, issue the **cd /opt/IBMsdd** and then **ls -l** commands to list all the installed SDD components. You will see output similar to the following:

```
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     88817 Sep 26 17:40 sdd-mod.o-2.4.2-smp
-rw-r----- 1 root    root     88689 Sep 26 17:40 sdd-mod.o-2.4.6-smp
-rw-r----- 1 root    root     89370 Sep 26 17:40 sdd-mod.o-2.4.9-smp
```

SDD utilities are packaged as executable files and contained in the `/bin` directory. If you issue the `cd /opt/IBMsdd/bin` and then `ls -l` commands, you will see output similar to the following:

```
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
```

Note: The `addpaths` command is still supported on the 2.4 kernels. On the 2.6 kernels `cfgvpath` will perform the functionality of `addpaths`.

If the installation failed, a message similar to the following is displayed:

```
package IBMsdd is not installed
```

Configuring SDD

Before you start the SDD configuration process, make sure that you have successfully configured the supported storage device to which your host system is attached and that the supported storage device is operational.

This section provides instructions for the following procedures:

- “Configuring and verifying an SDD” on page 209
- “Configuring SDD at system startup” on page 214
- “Maintaining SDD vpath device configuration persistence” on page 215

Table 21 lists all of the commands that can help system administrators configure SDD. More details about the function and use of each command are described later in this section.

Table 21. Summary of SDD commands for a Linux host system

Command	Description
<code>cfgvpath</code>	Configures SDD vpath devices. See the note.
<code>cfgvpath query</code>	Displays all SCSI disk devices.
<code>lsvpcfg</code>	Displays the current devices that are configured and their corresponding paths.
<code>rmvpath</code>	Removes one or all SDD vpath devices.

Table 21. Summary of SDD commands for a Linux host system (continued)

Command	Description
addpaths	<p>Adds any new paths to an existing SDD vpath device.</p> <p>This command is only supported for Linux 2.4 kernels.</p> <p>For Linux 2.6 kernels, the functionality of the addpaths command has been added to the cfgvpath command. If you need to add paths to an existing SDD vpath device with a Linux 2.6 kernel, run the cfgvpath command.</p>
sdd start	Loads the SDD driver and automatically configures disk devices for multipath access.
sdd stop	Unloads the SDD driver (requires that no SDD vpath devices currently be in use).
sdd restart	Unloads the SDD driver (requires that no SDD vpath devices currently be in use), and then loads the SDD driver and automatically configures disk devices for multipath access.

Note: For Linux 2.4 kernels, the SDD vpath devices are assigned names according to the following scheme:

```
vpatha, vpathb, ..., vpathp
vpathaa, vpathab, ..., vpathap
vpathba, vpathbb, ..., vpathbp
...
vpathza, vpathzb, ..., vpathzp
vpathaaa, vpathaab, ..., vpathaap
...
/
```

For Linux 2.6 kernels, the SDD vpath devices are assigned names according to the following scheme:

```
vpatha, vpathb, ..., vpathy, vpathz
vpathaa, vpathab, ..., vpathay, vpathaz
vpathba, vpathbb, ..., vpathby, vpathbz
...
vpathza, vpathzb, ..., vpathzy, vpathzz
vpathaaa, vpathaab, ..., vpathaay, vpathaaz
...
```

Configuring and verifying an SDD

Complete the following steps to load and configure SDD on your Linux host system:

SDD configuration

Use the following steps to configure SDD on your Linux host system.

1. Log on to your Linux host system as the root user.
2. Enter **sdd start**.
3. You can verify the configuration using the **datapath query device** command to determine that all your disk are configured. If the system is not configured properly, see “Verifying SDD configuration” on page 210.
4. Use the **sdd stop** command to unconfigure and unload the SDD driver. Use the **sdd restart** command to unconfigure, unload, and then restart the SDD

configuration process. If a vpath device is in use (mounted), then the **sdd stop** command fails with an error stating that the module `sdd-mod` is in use.

Verifying SDD configuration

Use the following steps to verify SDD configuration after running the **sdd start** command.

Note: If you are on an unsupported kernel you will get an error message about the kernel not being supported.

1. Enter **lsmod** or enter **cat /proc/modules** to verify that the SDD `sdd-mod` driver* is loaded. If it is successfully loaded, output similar to the following is displayed:

```
sdd-mod          233360  0 (unused)
qla2300          192000  0 (autoclean)
nls_iso8859-1    2880     1 (autoclean)
cs4232           3760     1 (autoclean)
ad1848           16752    0 (autoclean) [cs4232]
uart401          6352     0 (autoclean) [cs4232]
sound            56192    1 (autoclean) [cs4232 ad1848 uart401]
soundcore        4048     4 (autoclean) [sound]
nfsd              67664    4 (autoclean)
usb-uhci         20928    0 (unused)
usbcore          48320    1 [usb-uhci]
ipv6             131872   -1 (autoclean)
olympic          15856    1 (autoclean)
ipchains          34112    0 (unused)
lvm-mod          40880    0 (autoclean)
```

* For Linux 2.6 kernels, the SDD driver is displayed as `sdd_mod`.

2. Enter **cat /proc/IBMsdd** to verify that the SDD `sdd-mod` driver level matches that of your system kernel.

The following example shows that SDD 1.6.0.0 is installed on a Linux host system running a 2.4.9 symmetric multiprocessor kernel:

```
sdd-mod: SDD 1.6.0.0 2.4.9 SMP Sep 26 2001 17:39:06 (C) IBM Corp.
```

3. The order of disk recognition on a Linux system is:
 - a. Fibre-channel Host Bus Adapter (HBA) driver

The HBA driver needs to recognize the disks. The recognized disks are typically listed in `/proc/scsi/adapter_type/host_number`, for example `/proc/scsi/qla2300/2`. Example `/proc/scsi/adapter_type/host_number` output is shown below. Note that this is not always true for the Linux 2.6 kernel because the HBA driver version can use the `sysfs` filesystem instead of the `proc` filesystem to expose information.

- b. SCSI driver (`scsi-mod` or `scsi_mod`)

The SCSI driver has to recognize the disks, and, if this succeeds, it puts disk entries into `/proc/scsi/scsi`.

- c. SCSI disk driver (`sd-mod` or `sd_mod`)

The SCSI disk driver has to recognize the disk entries, and if this succeeds it puts the entries into `/proc/partitions`.

- d. SDD driver (`sdd-mod` or `sdd_mod`)

SDD then uses the disk entries in `/proc/partitions` to configure the SDD vpath devices. If configuration succeeds, SDD generates more entries in `/proc/partitions`.

Enter **cat /proc/scsi/adapter_type/N** to display the status of a specific adapter and the names of the attached devices. In this command, *adapter_type* indicates

the type of adapter that you are using, and *N* represents the host-assigned adapter number. The following example shows a sample output:

```
# ls /proc/scsi/
qla2300 scsi sym53c8xx
# ls /proc/scsi/qla2300/
2 3 HbaApiNode
# cat /proc/scsi/qla2300/2
QLogic PCI to Fibre Channel Host Adapter for ISP23xx:
    Firmware version: 3.01.18, Driver version 6.05.00b5
Entry address = e08ea060
HBA: QLA2300 , Serial# C81675
Request Queue = 0x518000, Response Queue = 0xc40000
Request Queue count= 128, Response Queue count= 512
Total number of active commands = 0
Total number of interrupts = 7503
Total number of IOCBs (used/max) = (0/600)
Total number of queued commands = 0
    Device queue depth = 0x10
Number of free request entries = 57
Number of mailbox timeouts = 0
Number of ISP aborts = 0
Number of loop resyncs = 47
Number of retries for empty slots = 0
Number of reqs in pending_q= 0, retry_q= 0, done_q= 0, scsi_retry_q= 0
Host adapter:loop state= <READY>, flags= 0x8a0813
Dpc flags = 0x0
MBX flags = 0x0
SRB Free Count = 4096
Port down retry = 008
Login retry count = 008
Commands retried with dropped frame(s) = 0

SCSI Device Information:
scsi-qla0-adapter-node=200000e08b044b4c;
scsi-qla0-adapter-port=210000e08b044b4c;
scsi-qla0-target-0=5005076300c70fad;
scsi-qla0-target-1=10000000c92113e5;
scsi-qla0-target-2=5005076300ce9b0a;
scsi-qla0-target-3=5005076300ca9b0a;
scsi-qla0-target-4=5005076801400153;
scsi-qla0-target-5=500507680140011a;
scsi-qla0-target-6=500507680140017c;
scsi-qla0-target-7=5005076801400150;
scsi-qla0-target-8=5005076801200153;
scsi-qla0-target-9=500507680120011a;
scsi-qla0-target-10=500507680120017c;
scsi-qla0-target-11=5005076801200150;

SCSI LUN Information:
(Id:Lun)
( 2: 0): Total reqs 35, Pending reqs 0, flags 0x0, 0:0:8c,
( 2: 1): Total reqs 29, Pending reqs 0, flags 0x0, 0:0:8c,
( 2: 2): Total reqs 29, Pending reqs 0, flags 0x0, 0:0:8c,
( 2: 3): Total reqs 29, Pending reqs 0, flags 0x0, 0:0:8c,
( 2: 4): Total reqs 29, Pending reqs 0, flags 0x0, 0:0:8c,
( 2: 5): Total reqs 29, Pending reqs 0, flags 0x0, 0:0:8c,
( 2: 6): Total reqs 29, Pending reqs 0, flags 0x0, 0:0:8c,
( 2: 7): Total reqs 29, Pending reqs 0, flags 0x0, 0:0:8c,
:
:
```

The disks that the QLogic adapter recognizes are listed at the end of the output under the heading **SCSI LUN Information**. The disk descriptions are shown one per line. An * at the end of a disk description indicates that the disk is not yet registered with the operating system. SDD cannot configure devices that are not registered with the operating system. See the appropriate *Host Systems Attachment Guide* for your product to learn about SCSI LUN discovery in Linux.

4. Enter `cfgvpath query` to verify that you have configured the SCSI disk devices that you allocated and configured for SDD. The `cfgvpath query` is effectively looking at the `/proc/partitions` output.

After you enter the `cfgvpath query` command, a message similar to the following is displayed. This example output is for a system with disk storage system and virtualization product LUNs.

```
/dev/sda ( 8,  0) host=0 ch=0 id=0 lun=0 vid=IBM      pid=DDYS-T36950M serial=xxxxxxxxxxxx      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0 X
/dev/sdb ( 8, 16) host=2 ch=0 id=0 lun=0 vid=IBM      pid=2105E20   serial=60812028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdc ( 8, 32) host=2 ch=0 id=0 lun=1 vid=IBM      pid=2105E20   serial=70912028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdd ( 8, 48) host=2 ch=0 id=0 lun=2 vid=IBM      pid=2105E20   serial=31812028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sde ( 8, 64) host=2 ch=0 id=0 lun=3 vid=IBM      pid=2105E20   serial=31C12028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdf ( 8, 80) host=2 ch=0 id=1 lun=0 vid=IBM      pid=2105E20   serial=60812028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdg ( 8, 96) host=2 ch=0 id=1 lun=1 vid=IBM      pid=2105E20   serial=70912028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdh ( 8, 112) host=2 ch=0 id=1 lun=2 vid=IBM      pid=2105E20   serial=31812028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdi ( 8, 128) host=2 ch=0 id=1 lun=3 vid=IBM      pid=2105E20   serial=31C12028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdj ( 8, 144) host=2 ch=0 id=6 lun=0 vid=IBM      pid=2145      serial=600507680183000a80000000000000a ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdk ( 8, 160) host=2 ch=0 id=6 lun=1 vid=IBM      pid=2145      serial=600507680183000a80000000000000b ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdl ( 8, 176) host=2 ch=0 id=6 lun=2 vid=IBM      pid=2145      serial=600507680183000a80000000000000c ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdm ( 8, 192) host=2 ch=0 id=6 lun=3 vid=IBM      pid=2145      serial=600507680183000a80000000000000d ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdn ( 8, 208) host=2 ch=0 id=6 lun=4 vid=IBM      pid=2145      serial=600507680183000a80000000000000e ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdo ( 8, 224) host=2 ch=0 id=6 lun=5 vid=IBM      pid=2145      serial=600507680183000a80000000000000f ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdp ( 8, 240) host=2 ch=0 id=6 lun=6 vid=IBM      pid=2145      serial=600507680183000a800000000000010 ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdq ( 65,  0) host=2 ch=0 id=6 lun=7 vid=IBM      pid=2145      serial=600507680183000a800000000000011 ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdr ( 65, 16) host=2 ch=0 id=6 lun=8 vid=IBM      pid=2145      serial=600507680183000a800000000000012 ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sds ( 65, 32) host=2 ch=0 id=6 lun=9 vid=IBM      pid=2145      serial=600507680183000a800000000000013 ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdt ( 65, 48) host=2 ch=0 id=7 lun=0 vid=IBM      pid=2145      serial=600507680183000a80000000000000a ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdu ( 65, 64) host=2 ch=0 id=7 lun=1 vid=IBM      pid=2145      serial=600507680183000a80000000000000b ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdv ( 65, 80) host=2 ch=0 id=7 lun=2 vid=IBM      pid=2145      serial=600507680183000a80000000000000c ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdw ( 65, 96) host=2 ch=0 id=7 lun=3 vid=IBM      pid=2145      serial=600507680183000a80000000000000d ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdx ( 65, 112) host=2 ch=0 id=7 lun=4 vid=IBM      pid=2145      serial=600507680183000a80000000000000e ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdy ( 65, 128) host=2 ch=0 id=7 lun=5 vid=IBM      pid=2145      serial=600507680183000a80000000000000f ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdz ( 65, 144) host=2 ch=0 id=7 lun=6 vid=IBM      pid=2145      serial=600507680183000a800000000000010 ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdaa ( 65, 160) host=2 ch=0 id=7 lun=7 vid=IBM      pid=2145      serial=600507680183000a800000000000011 ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdab ( 65, 176) host=2 ch=0 id=7 lun=8 vid=IBM      pid=2145      serial=600507680183000a800000000000012 ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdac ( 65, 192) host=2 ch=0 id=7 lun=9 vid=IBM      pid=2145      serial=600507680183000a800000000000013 ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdad ( 65, 208) host=2 ch=0 id=10 lun=0 vid=IBM      pid=2145      serial=600507680183000a80000000000000a ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdae ( 65, 224) host=2 ch=0 id=10 lun=1 vid=IBM      pid=2145      serial=600507680183000a80000000000000b ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdaf ( 65, 240) host=2 ch=0 id=10 lun=2 vid=IBM      pid=2145      serial=600507680183000a80000000000000c ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
:
```

```
/dev/sda ( 8,  0) host=0 ch=0 id=0 lun=0 vid=IBM      pid=DDYS-T36950M serial=xxxxxxxxxxxx      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0 X
/dev/sdb ( 8, 16) host=2 ch=0 id=0 lun=0 vid=IBM      pid=2105E20   serial=60812028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdc ( 8, 32) host=2 ch=0 id=0 lun=1 vid=IBM      pid=2105E20   serial=70912028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdd ( 8, 48) host=2 ch=0 id=0 lun=2 vid=IBM      pid=2105E20   serial=31812028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sde ( 8, 64) host=2 ch=0 id=0 lun=3 vid=IBM      pid=2105E20   serial=31C12028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdf ( 8, 80) host=2 ch=0 id=1 lun=0 vid=IBM      pid=2105E20   serial=60812028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdg ( 8, 96) host=2 ch=0 id=1 lun=1 vid=IBM      pid=2105E20   serial=70912028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdh ( 8, 112) host=2 ch=0 id=1 lun=2 vid=IBM      pid=2105E20   serial=31812028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdi ( 8, 128) host=2 ch=0 id=1 lun=3 vid=IBM      pid=2105E20   serial=31C12028      ctrl_flag=0 ctrl_nbr=0 df_ctlr=0
/dev/sdj ( 8, 144) host=2 ch=0 id=6 lun=0 vid=IBM      pid=2145      serial=600507680183000a80000000000000a ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdk ( 8, 160) host=2 ch=0 id=6 lun=1 vid=IBM      pid=2145      serial=600507680183000a80000000000000b ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdl ( 8, 176) host=2 ch=0 id=6 lun=2 vid=IBM      pid=2145      serial=600507680183000a80000000000000c ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdm ( 8, 192) host=2 ch=0 id=6 lun=3 vid=IBM      pid=2145      serial=600507680183000a80000000000000d ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdn ( 8, 208) host=2 ch=0 id=6 lun=4 vid=IBM      pid=2145      serial=600507680183000a80000000000000e ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdo ( 8, 224) host=2 ch=0 id=6 lun=5 vid=IBM      pid=2145      serial=600507680183000a80000000000000f ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdp ( 8, 240) host=2 ch=0 id=6 lun=6 vid=IBM      pid=2145      serial=600507680183000a800000000000010 ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdq ( 65,  0) host=2 ch=0 id=6 lun=7 vid=IBM      pid=2145      serial=600507680183000a800000000000011 ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdr ( 65, 16) host=2 ch=0 id=6 lun=8 vid=IBM      pid=2145      serial=600507680183000a800000000000012 ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sds ( 65, 32) host=2 ch=0 id=6 lun=9 vid=IBM      pid=2145      serial=600507680183000a800000000000013 ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdt ( 65, 48) host=2 ch=0 id=7 lun=0 vid=IBM      pid=2145      serial=600507680183000a80000000000000a ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdu ( 65, 64) host=2 ch=0 id=7 lun=1 vid=IBM      pid=2145      serial=600507680183000a80000000000000b ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdv ( 65, 80) host=2 ch=0 id=7 lun=2 vid=IBM      pid=2145      serial=600507680183000a80000000000000c ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdw ( 65, 96) host=2 ch=0 id=7 lun=3 vid=IBM      pid=2145      serial=600507680183000a80000000000000d ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdx ( 65, 112) host=2 ch=0 id=7 lun=4 vid=IBM      pid=2145      serial=600507680183000a80000000000000e ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdy ( 65, 128) host=2 ch=0 id=7 lun=5 vid=IBM      pid=2145      serial=600507680183000a80000000000000f ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdz ( 65, 144) host=2 ch=0 id=7 lun=6 vid=IBM      pid=2145      serial=600507680183000a800000000000010 ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdaa ( 65, 160) host=2 ch=0 id=7 lun=7 vid=IBM      pid=2145      serial=600507680183000a800000000000011 ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdab ( 65, 176) host=2 ch=0 id=7 lun=8 vid=IBM      pid=2145      serial=600507680183000a800000000000012 ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdac ( 65, 192) host=2 ch=0 id=7 lun=9 vid=IBM      pid=2145      serial=600507680183000a800000000000013 ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdad ( 65, 208) host=2 ch=0 id=10 lun=0 vid=IBM      pid=2062      serial=600507680183000a80000000000000a ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
/dev/sdae ( 65, 224) host=2 ch=0 id=10 lun=1 vid=IBM      pid=2062      serial=600507680183000a80000000000000b ctrl_flag=1 ctrl_nbr=1 df_ctlr=0
/dev/sdaf ( 65, 240) host=2 ch=0 id=10 lun=2 vid=IBM      pid=2062      serial=600507680183000a80000000000000c ctrl_flag=1 ctrl_nbr=0 df_ctlr=0
:
```

The sample output shows the name and serial number of the SCSI disk device, its connection information, and its product identification. A capital letter X at the end of a line indicates that SDD currently does not support the device or the device is in use and `cfgvpath` has not configured it. The `cfgvpath` utility examines `/etc/fstab` and the output of the `mount` command in order to determine the disks that it should not configure. If `cfgvpath` has not configured a disk that you think it should have configured, verify that an entry for one of these disks exists in `/etc/fstab` or in the output of the `mount` command. If the entry is incorrect, delete the wrong entry and issue `cfgvpath` again to configure the device.

Automatic SCSI disk device exclusion during SDD configuration

The SDD configuration might sometimes exclude a SCSI disk (sd) device that is present on the system from being configured for use by an SDD vpath device in the following situations:

1. The SCSI disk device is from an unsupported storage subsystem.

You can determine whether your SCSI disk devices are supported by running **cfgvpath query** and checking the output. See “Configuring and verifying an SDD” on page 209 for additional information about how to determine whether the SCSI disk devices are supported.

2. The SCSI disk device is listed in the file `/etc/fstab`.
fstab is a configuration file that contains information about the important file system information regarding disk devices and partitions, such as how and where they should be mounted. For example, an entry specifying the disk or partition that acts as swap space would be in fstab. The system administrator must keep the fstab configuration file up-to-date so that when SDD checks this file, it is able to correctly exclude drives and partitions.
3. The SCSI disk device is currently mounted (using the Linux **mount** command). SDD configuration assumes that the device is in use for another purpose and will not configure the device.
4. The SCSI disk device is currently bound to a raw device. Use the **raw -qa** command to check the raw device bindings. If the major, minor pair of the raw command output matches with an SCSI disk device major, minor pair, the SCSI disk device will be excluded.

Important things to note about the exclusion process are:

1. When running **cfgvpath** or **sdd start**, the SDD configuration will print out a message indicating whether it has excluded any SCSI disk devices.
2. Once an SCSI disk device that belongs to an SDD vpath device is excluded, all SCSI disk devices (or paths) belonging to the SDD vpath device will be excluded.

Reconfiguring using SDD userspace commands

You can use the following commands when reconfiguring SDD vpath devices:

cfgvpath

Enter **cfgvpath** to configure SDD vpath devices.

The configuration information is saved by default in the `/etc/vpath.conf` file to maintain vpath name persistence in subsequent driver loads and configurations. You might choose to specify your own configuration file by issuing the **cfgvpath -f configuration_file_name.cfg** command where *configuration_file_name* is the name of the configuration file that you want to specify. If you use a self-specified configuration file, SDD only configures the SDD vpath devices that this file defines.

Enter **cfgvpath ?** for more information about the **cfgvpath** command.

rmvpath

You can remove an SDD vpath device by using the **rmvpath vpath_name** command, where *vpath_name* represents the name of the SDD vpath device that is selected for removal.

Enter **rmvpath ?** for more information about the **rmvpath** command.

lsvpcfg

Verify the SDD vpath device configuration by entering **lsvpcfg** or **datapath query device**.

If you successfully configured SDD vpath devices, output similar to the following is displayed by **lsvpcfg**. This example output is for a system with disk storage system and virtualization product LUNs:

```
sdd-mod: SDD 1.6.0.0 2.4.19-64GB-SMP SMP Mar 3 2003 18:06:49 (C) IBM Corp.
000 vpatha ( 247, 0) 60812028 = /dev/sdb /dev/sdf /dev/sdax /dev/sdbb
001 vpathb ( 247, 16) 70912028 = /dev/sdc /dev/sdg /dev/sday /dev/sdbc
002 vpathc ( 247, 32) 31B12028 = /dev/sdd /dev/sdh /dev/sdaz /dev/sdbd
003 vpathd ( 247, 48) 31C12028 = /dev/sde /dev/sdi /dev/sdba /dev/sdbe
004 vpathe ( 247, 64) 600507680183000a800000000000000a = /dev/sdj /dev/sdt /dev/sdad /dev/sdan /dev/sdbf /dev/sdbp /dev/sdbz /dev/sdcj
005 vpathf ( 247, 80) 600507680183000a800000000000000b = /dev/sdk /dev/sdu /dev/sdae /dev/sdao /dev/sdbg /dev/sdbq /dev/sdca /dev/sdck
006 vpathg ( 247, 96) 600507680183000a800000000000000c = /dev/sdl /dev/sdv /dev/sdaf /dev/sdap /dev/sdbh /dev/sdbx /dev/sdcb /dev/sdcl
007 vpathh ( 247, 112) 600507680183000a800000000000000d = /dev/sdm /dev/sdw /dev/sdag /dev/sdaq /dev/sdbi /dev/sdbs /dev/sdcc /dev/sdcm
008 vpathi ( 247, 128) 600507680183000a800000000000000e = /dev/sdn /dev/sdx /dev/sdah /dev/sdar /dev/sdbj /dev/sdbt /dev/sdcd /dev/sdcn
009 vpathj ( 247, 144) 600507680183000a800000000000000f = /dev/sdo /dev/sdy /dev/sdai /dev/sdas /dev/sdbk /dev/sdbu /dev/sdce /dev/sdco
010 vpathk ( 247, 160) 600507680183000a8000000000000010 = /dev/sdp /dev/sdz /dev/sdaj /dev/sdat /dev/sdbl /dev/sdbv /dev/sdcf /dev/sdcp
011 vpathl ( 247, 176) 600507680183000a8000000000000011 = /dev/sdq /dev/sdaa /dev/sdak /dev/sdau /dev/sdbm /dev/sdbw /dev/sdcg /dev/sdcq
012 vpathm ( 247, 192) 600507680183000a8000000000000012 = /dev/sdr /dev/sdab /dev/sdal /dev/sdav /dev/sdbn /dev/sdbx /dev/sdch /dev/sdcr
013 vpathn ( 247, 208) 600507680183000a8000000000000013 = /dev/sds /dev/sdac /dev/sdam /dev/sdaw /dev/sdbo /dev/sdbz /dev/sdc1 /dev/sdcs
```

See Chapter 10, “Using the datapath commands,” on page 371 for more information about the **datapath query device** command and all other SDD datapath commands.

addpaths

You can issue the **addpaths** command to add paths to existing SDD vpath devices. For SDD to discover new paths, the Linux kernel SCSI disk driver must already be aware of the path. For example, **addpaths** would be useful in a scenario where disks are configured and are visible to the OS but were unavailable at the time that SDD was configured because of a failed switch or unplugged fibre cable. Later, when the disks are recovered through the recovery process or maintenance, you can issue **addpaths** on a running system to add the restored paths.

Issue the **addpaths** command to add new paths to existing disks. Use **cfgvpath** to add new disks. See “Using dynamic reconfiguration” on page 217.

Note: For Linux 2.6 kernels, **addpaths** is not supported because the functionality of **addpaths** has been incorporated into the **cfgvpath** command. To add new paths to existing disks when using a Linux 2.6 kernel, run **cfgvpath**.

Configuring SDD at system startup

Note: SDD is currently set to not be loaded on system startup after installation. Use this section to load SDD on system startup. A rpm upgrade does not change the current configuration.

SDD can be set to automatically load and configure when your Linux system starts up. SDD provides a startup script `sdd.rcscript` file in the `/opt/IBMsdd/bin` directory and creates a symbolic link to `/etc/init.d/sdd`.

Complete the following steps to configure SDD at system startup:

1. Log on to your Linux host system as the root user.
2. At startup run one of these commands to enable run level X:

For Red Hat:

```
chkconfig --level X sdd on
```

For SUSE:

```
chkconfig --set sdd X
```

3. Enter **chkconfig --list sdd** to verify that the system startup option is enabled for SDD configuration.
4. Restart your host system so that SDD is loaded and configured.

If necessary, you can disable the startup option by entering:

```
chkconfig --level X sdd off
```

In order for SDD to automatically load and configure, the host bus adapter (HBA) driver must already be loaded. This can be assured at start time by adding the appropriate driver or drivers to the kernel's initial RAM disk. See the Red Hat **mkinitrd** command documentation or the SUSE **mk_initrd** command documentation for more information. Additional suggestions might be available from the HBA driver vendor.

Sometimes certain system configurations require SDD to start earlier than is possible under the procedure described above. The general rule is: if some application, filesystem, or other product needs to use an SDD vpath device before it is loaded in the system init scripts, then you will need to use another procedure to start SDD to allow these applications or filesystems access to SDD vpath devices. Some of the known system configurations are described below. This is not an exhaustive list, but it does provide an idea of situations where other methods are required:

1. SDD remote boot

If booting off of a SDD vpath device, SDD needs to be available before the root filesystem is mounted. This means SDD needs to be placed in the initial ramdisk (initrd). See “Bootting Linux over the SAN with SDD” on page 224 for more instructions on how to set up this environment.

2. Linux Logical Volume Manager (LVM) with SDD

Linux LVM with SDD often requires SDD to start early in the init script process because the LVM initialization occurs relatively early. If LVM is used to encapsulate the root disk, SDD needs to be placed in the initial ramdisk (initrd). See “Using LVM with SDD” on page 78 for more information.

Any other customized application, driver, or filesystem that requires access to a SDD vpath device early in the boot process might require: (1) SDD be placed in the initial ramdisk (initrd), or (2) the SDD startup script be placed in earlier in the init scripts.

Maintaining SDD vpath device configuration persistence

Use the **cfgvpath** command to configure SDD vpath devices. For first time configuration, the configuration method finds all SCSI disk devices, then configures and assigns SDD vpath devices accordingly. The configuration is saved in `/etc/vpath.conf` to maintain name persistence in subsequent driver loads and configurations.

The `/etc/vpath.conf` is not modified during a rpm upgrade (**rpm -U**). However, if the rpm is removed and reinstalled (using the **rpm -e** and **rpm -i** commands), the `/etc/vpath.conf` is removed. If you are doing a rpm removal, it is important to manually save your `/etc/vpath.conf` and restore it after the rpm has been reinstalled, before issuing **sdd start**.

After the SDD vpath devices are configured, issue **lsvpcfg** or the **datapath query device** command to verify the configuration. See “datapath query device” on page 378 for more information.

You can manually exclude a device in `/etc/vpath.conf` from being configured. To manually exclude a device from being configured, edit the `vpath.conf` file prior to

running `sdd start`, adding a `#` before the first character of the entry for the device that you want to remain unconfigured. Removing the `#` allows a previously excluded device to be configured again.

The following output shows the contents of a `vpath.conf` file with `vpathb` and `vpathh` not configured:

```
vpatha 60920530
#vpathb 60A20530
vpathc 60B20530
vpathd 60C20530
vpathe 70920530
vpathf 70A20530
vpathg 70B20530
#vpathh 70C20530
```

Dynamically changing the SDD path-selection policy algorithm

SDD supports path-selection policies that increase the performance of multipath-configured supported storage devices and make path failures transparent to applications. The following path-selection policies are supported:

failover only (fo)

All I/O operations for the device are sent to the same (preferred) path until the path fails because of I/O errors. Then an alternate path is chosen for subsequent I/O operations.

load balancing (lb)

The path to use for an I/O operation is chosen by estimating the load on the adapter to which each path is attached. The load is a function of the number of I/O operations currently in process. If multiple paths have the same load, a path is chosen at random from those paths. Load-balancing mode also incorporates failover protection.

The load-balancing policy is also known as the *optimized policy*.

load balancing sequential (lbs)

This policy is the same as the load-balancing policy with optimization for sequential I/O.

The load-balancing sequential policy is also known as the *optimized sequential policy*. This is the default setting.

round robin (rr)

The path to use for each I/O operation is chosen at random from paths that were not used for the last I/O operation. If a device has only two paths, SDD alternates between the two.

round robin sequential (rrs)

This policy is the same as the round-robin policy with optimization for sequential I/O.

The default path-selection policy for an SDD device is load balancing sequential. You can change the policy for an SDD device. SDD supports dynamic changing of the SDD devices' path-selection policy.

Before changing the path-selection policy, determine the active policy for the device. Enter **datapath query device N** where *N* is the device number of the SDD vpath device to show the current active policy for that device.

datapath set device policy command

Use the **datapath set device policy** command to change the SDD path-selection policy dynamically:

See “datapath set device policy” on page 389 for more information about the **datapath set device policy** command.

Using dynamic reconfiguration

Before SDD can configure a dynamically-added LUN, the OS and HBA drivers must detect and configure it. Typically, to see a LUN that has been dynamically added, a reload of the HBA driver is required. However, specific versions of HBA drivers and adapters might offer procedures to force a re-scan of the bus in order to detect new LUNs without requiring a driver reload. In general, the QLogic and Emulex drivers for Linux 2.6 kernels (that is, SLES 9 and RHEL 4) offer procedures using the sysfs interface. The HBA drivers for the Linux 2.4 kernel have more limited options. See the respective HBA vendor's documentation to determine if a procedure is supported and, if so, with which HBA driver versions.

If a HBA driver reload is used, it is important to run **sdd stop** before unloading the HBA driver and **sdd start** after the HBA driver has been reloaded. If a procedure is used that does not require a HBA reload, running **cfgvpath** will be adequate to configure the added LUN. See the note below.

The **cfgvpath** command finds the current hardware configuration and compares it to the SDD vpath device configuration in memory and then identifies a list of differences. It then issues commands to update the SDD vpath device configuration in memory with the current hardware configuration. The commands that **cfgvpath** issues to the SDD driver are:

- Add an SDD vpath device.
- Remove an SDD vpath device; this will fail if device is busy.
- Add path to an SDD vpath device. (Only for SDD drivers on Linux 2.6 kernels.)
- Remove path for an SDD vpath device; this will fail deletion of path if device is busy, but will set path to DEAD and OFFLINE. (Only for SDD drivers on Linux 2.6 kernels.)

Note: Although unlikely, it is possible that the procedure provided by the HBA vendor to force a re-scan of the bus can cause SCSI device name slippage in the OS, which means the same device file name refers to different LUNs before and after the re-scan. For example, before the dynamic LUN addition, `/dev/sda` denoted a path to LUN 0, but after LUN 1 is added dynamically, `/dev/sda` now refers to a path to LUN 1. In such a scenario, running **sdd restart** is required. In SDD 1.6.1.0 versions and later, **cfgvpath** will automatically refuse to configure the added LUN if name slippage has occurred.

Uninstalling SDD

You must unload the SDD driver before uninstalling SDD. Complete the following steps to remove SDD from a Linux host system:

1. Log on to your Linux host system as the root user.
2. Enter **sdd stop** to remove the driver.
3. Enter **rpm -e IBMsd** to remove the SDD package.
4. Verify the SDD removal by entering either **rpm -q IBMsd** or **rpm -ql IBMsd**.

If you successfully removed SDD, output similar to the following is displayed:

```
package IBMsdd is not installed
```

Note: The `sdd stop` command will not unload a driver that is in use.

Mounting SDD devices with automount

This procedure describes how to automatically mount vpath partitions using the program **automount**. **automount** is a daemon that includes a kernel component that will automatically mount partitions when accessed.

For example, if you had an SDD device `vpatha` configured to automount at the mount point `/mnt/vpatha`, `vpatha` would not be mounted until any program tried to access something on that mount point. Once **automount** has automatically mounted the device, it can also automatically unmount the device after a certain timeout value, unless otherwise specified.

Setting up automount

The **autofs** daemon should be set up at boot time by default. To check this, issue the following command:

```
chkconfig --list autofs
```

The output of the command should state the runlevels to which **autofs** is set. For example:

```
autofs          0:off  1:off  2:off  3:on   4:on   5:on   6:off
```

This output indicates that **autofs** is running on runlevels 3, 4, and 5, which should be the default setting. If you notice that the **autofs** daemon is not running on runlevels 3, 4, 5, issue the following commands to ensure that it will run on startup:

On SUSE:

```
chkconfig autofs 345
```

On Red Hat:

```
chkconfig --level 345 autofs on
```

Configuring automount

Use the following steps to configure automount:

1. Configure the master map file.

Automount configuration requires the configuration of the master map file, `/etc/auto.master`.

The format of the file is the following:

```
[mount point] [map file] [options]
```

where,

mount point

This variable will be the master mount point under which all the vpath devices will be mounted. For example, /mnt, or /vpath (note that it is an absolute path).

Note: The mount point that you specify will be mounted over by **autofs**. That means that whatever items you had mounted at that mount point will be invisible once **automount** is activated. Thus, ensure that you do not have conflicting mount points for separate applications and that if you plan to mount other things under the master mount point, you do so with **automount** and not within **fstab** or another facility or script. If the conflict is unavoidable, change the automount master mount point to a nonconflicting mount point to prevent problems from occurring.

map file

This is another separate file that will describe under which names certain devices will be mounted and the mount variables for the device. Usually, it is named after the mount point, such as auto.mnt or auto.vpath. It will usually reside under /etc.

options These are the options that you can specify and can be referenced by the automount man page.

The most relevant setting is the **--timeout** setting. The timeout setting is the number of seconds that automount will wait for mount point access before unmounting that mount point. If you set this value to 0, automount will not attempt to unmount the master mount point (that is, it will remain permanently mounted unless it is manually unmounted). The default setting is 5 minutes.

The following example shows a sample auto.master file:

```
/vpath /etc/auto.vpath --timeout=0
```

2. Configure the secondary map file.

The secondary map file is the file referred to by the file /etc/auto.master. The format of this map file is:

```
[secondary mount point] [mount options] [device name]
```

where,

secondary mount point

The secondary mount point is the mount point relative to the master mount point. For example, if you wanted vpatha to be mounted at /vpath/vpatha, you would set this secondary mount point to *vpatha*.

mount options

The mount options are standard options passed to the Linux mount command using the *-o* option. The only difference is that you can use the option *fstype* to specify the exact filesystem type of the device. For example, you can use *ext2*, *ext3*, *reiserfs*, etc for the *fstype*. You can find the other options under the man page for mount.

Set the *fstype* to the correct value and use the two options *defaults* and *check=normal*. Defaults will give some values to the filesystem that are standard for most Linux operating environments. The *check=normal* option ensures that certain sanity checks are made on the filesystem before mounting. You can set *check=strict* to ensure even stricter

checking rules during mount time; but performance might be degraded. Most modern filesystems check themselves after a certain number of mounts.

device name

The following example shows a sample `auto.vpath` file:

```
vpatha -fstype=ext3,defaults,check=normal :/dev/vpatha
vpathi -fstype=ext2,defaults,check=normal :/dev/vpathi
```

3. Capture your new file settings. Test with a reboot cycle at least once to ensure that `autofs` is loaded with the current map files and that the system will automatically mount the devices correctly. Perform one of the following steps:
 - Reboot the system.
 - Run `/etc/init.d/autofs restart`.

Sample mount output

The following samples shows example output.

Mount output with `automount` loaded and before accessing `/vpath/vpatha`:

```
/dev/hda3 on / type ext3 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
/dev/hda1 on /boot type ext3 (rw)
none on /dev/shm type tmpfs (rw)
automount(pid14846) on /vpath type autofs (rw,fd=4,pgrp=14846,minproto=2,maxproto=4)
```

Mount output after accessing `/vpath/vpatha`:

```
/dev/hda3 on / type ext3 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
/dev/hda1 on /boot type ext3 (rw)
none on /dev/shm type tmpfs (rw)
automount(pid14846) on /vpath type autofs (rw,fd=4,pgrp=14846,minproto=2,maxproto=4)
/dev/vpatha on /vpath/vpatha type ext3 (rw,check=normal)
```

Mount output after accessing `/vpath/vpathi`:

```
/dev/hda3 on / type ext3 (rw)
none on /proc type proc (rw)
none on /dev/pts type devpts (rw,gid=5,mode=620)
/dev/hda1 on /boot type ext3 (rw)
none on /dev/shm type tmpfs (rw)
automount(pid16309) on /vpath type autofs (rw,fd=4,pgrp=16309,minproto=2,maxproto=4)
/dev/vpatha on /vpath/vpatha type ext3 (rw,check=normal)
/dev/vpathi on /vpath/vpathi type ext2 (rw,check=normal)
```

Use **automount** to mount SDD `vpath` devices. However, on Red Hat Linux, if you want to add the mount points over SDD devices in `/etc/fstab` and have them mounted by **mount -a** during system startup time, you must not enable the `autofsck` option (which is done by setting the sixth field of the new entry to 0). Also, ensure that you make the following change in `/opt/IBMsdd/bin/sdd.rcscript`:

Change:

```
# chkconfig: - 25 75
```

to:

```
# chkconfig: - 00 75
```

This allows the SDD driver to start as early as possible so that other RC scripts that mount file systems as part of the startup sequence will mount vpath devices configured by SDD.

The usual startup script that calls `mount -a` is `S25netfs`. If this script is not enabled, either enable it using `chkconfig`, or add the `mount -a` command to `rc.local` so that any entries in `/etc/fstab` that have not yet been mounted will be attempted. Also, verify that other applications that need to use SDD vpath devices or their mount points are started after SDD has been started, configured, and its filesystems mounted.

You then need to issue `chkconfig sdd on` to configure SDD at system startup. Use `chkconfig --list sdd` to verify the run levels that `sdd` is configured to start. If the run levels are not correct, modify the run levels by using the `--level` option in `chkconfig` to adjust the levels to values that are appropriate for the system configuration.

Using Linux Logical Volume Manager with SDD

Using Linux Logical Volume Manager (LVM) with SDD requires some configuration to ensure the correct behavior. LVM 2 (on 2.6 kernels, that is, SLES 9 and RHEL 4) requires some additional configuration that is discussed at the end of this section.

In order to use SDD with the Linux LVM, the SDD module must be loaded and vpath devices configured before LVM is loaded. This is because LVM has a dependency on SDD vpath devices, and therefore SDD must be loaded prior to LVM startup.

Note: This procedure will describe how to set up SDD to load before LVM using boot up scripts. This procedure might not work if LVM is loaded in the `initrd` (initial ramdisk). One common reason why LVM would be loaded in the `initrd` is because the system is booting off of an LVM device for its root disk. Supporting this type of environment requires that SDD also be placed in the `initrd`, which is a process that is not described here.

The procedures and requirements are different for SUSE and Red Hat.

Using LVM on SUSE

LVM with SDD on SLES 8 is only supported on LVM version `lvm-1.0.5-129` or later. Upgrade your LVM to this version or later using the SUSE maintweb update. Your LVM will not work with SDD vpath devices correctly otherwise.

LVM 2 with SDD on SLES 9 will work after following the LVM 2-specific directions at the end of this section.

In order to get SDD to load before LVM, you must load SDD at boot time using the bootup scripts. Complete the following steps:

1. Enter `cd /etc/init.d/` to change to the `/etc/init.d/` directory.
2. Enter `cp /opt/IBMsdd/bin/sdd.rcscript boot.sdd` to copy the `boot.sdd` file to the `/etc/init.d/` directory.
3. Edit the `boot.sdd` file and add the following line to the beginning of the start function:

```
mount -n -o remount,rw / 2> /dev/null
```

Because the SDD configuration utility (**cfgvpath**) must be able to write certain configuration parameters to the root disk, this line is needed to remount the root filesystem in read/write mode.

4. Add another line at the end of the start function to remount the root filesystem back into read-only mode. This restores the mount state before and after the start function. The system remounts the filesystem to read write at a later point in the boot process. Use the following line to remount in read-only mode:

```
mount -n -o remount,ro / 2> /dev/null (the only change from above is ro)
```

The beginning of your start function should look like the following:

```
start() {
    mount -n -o remount,rw / 2> /dev/null # ADDED THIS LINE
    echo -n "Starting $dev_name driver load: "
    rm -f ${driver_dir}/${driver}.o
    ...
    mount -n -o remount,ro / 2> /dev/null # ADDED THIS LINE
}
```

5. Enter `cd /etc/init.d/boot.d` to change to the `/etc/init.d/boot.d` directory.
6. Create a link named `Sxxboot.sdd` with `xx` being a number smaller than the one that is on the LVM boot script link. For example, the LVM on this system is at `S04boot.lvm`:

```
# ls -l | grep lvm
lrwxrwxrwx 1 root root 11 Aug 12 17:35 S04boot.lvm -> ../boot.lvm*
```

Because LVM is loading up at `S04`, you must set SDD to at least `S03` in order to avoid this problem. Therefore, create a link to the `boot.sdd` file that was just modified:

```
# ln -s ../boot.sdd S03boot.sdd
```

Here are the results:

```
# ls -l
...
lrwxrwxrwx 1 root root 11 Mar 11 12:03 S03boot.sdd -> ../boot.sdd*
lrwxrwxrwx 1 root root 11 Aug 12 17:35 S04boot.lvm -> ../boot.lvm*
...
```

Because SUSE uses the numbering scheme to determine which script is run first at boot time, you must ensure that the SDD script is run before the LVM script is run.

7. If you have SDD starting in the runlevel init scripts, you must shut off the script. To do this, issue the **chkconfig** command:

```
chkconfig -s sdd off
```

8. Configure LVM. Reboot the system and the LVM configuration should come up after reboot using SDD vpath devices.

Using LVM on Red Hat

LVM with SDD on RHEL 3 is only supported on LVM version lvm-1.0.8-14 or later. This is the LVM level that comes packaged with RHEL 3 Update 6. Upgrade your LVM to this version or later using the Red Hat **up2date** utility. Your LVM will not work with SDD vpath devices correctly otherwise.

LVM 2 with SDD on RHEL 4 will work after following the LVM 2-specific directions at the end of this section.

In order to get SDD to load before LVM, use the following steps to cause SDD to load at boot time.

1. Edit the file `/etc/rc.sysinit`.
2. Locate the block of commands that start with the following:

```
# Remount the root filesystem read-write.
update_boot_stage RCmountfs
state=`awk '/ \/ / && ($3 !~ /rootfs/) { print $4 }' /proc/mounts`
[ "$state" != "rw" -a "$READONLY" != "yes" ] && \
action $"Remounting root filesystem in read-write mode: " mount -n -o remount,rw /

# LVM initialization
...
```

3. Append the following to the end of the block of commands, before the `# LVM` initialization comment, or on RHEL 4, before the `# LVM2` initialization comment:

```
# Starting SDD
/etc/init.d/sdd start
```

4. The affected section of the `rc.sysinit` file should look like this:

```
# Remount the root filesystem read-write.
update_boot_stage RCmountfs
state=`awk '/ \/ / && ($3 !~ /rootfs/) { print $4 }' /proc/mounts`
[ "$state" != "rw" -a "$READONLY" != "yes" ] && \
action $"Remounting root filesystem in read-write mode: " mount -n -o remount,rw /

# Starting SDD
/etc/init.d/sdd start

# LVM initialization
...
```

5. If you have SDD starting in the runlevel init scripts, you need to shut off the script. You can do this using the **chkconfig** command.

```
chkconfig sdd off
```

6. Configure LVM.
7. Reboot the system and the LVM configuration should come up after reboot using SDD vpath devices.

LVM 2 general configuration information

The information in this section only applies to systems running LVM version 2. This means that the system should be running a 2.6 kernel, such as SLES 9 or RHEL 4.

LVM2 refers to a new userspace toolset that provides logical volume management facilities on Linux. In order for LVM2 to work on your machine, you must have the LVM2 rpm installed; look on the OS install CD's for the LVM2 package.

In order to get LVM2 to run with SDD, several changes must be made to the LVM configuration file `/etc/lvm/lvm.conf`. In this file, there are several fields that must be modified. As always, back up the file before attempting to modify it.

Under the section **devices**, there are two values that must be modified.

1. Value of *filter*.

In the file, you might see the default as the following:

```
# By default, accept every block device except udev names:
filter = [ "r|/dev./*/by-path./*", "r|/dev./*/by-id./*", "a./*/" ]
```

This filter is too broad for SDD, because it recognizes both SDD vpath devices as well as the underlying paths (that is, `/dev/sdxxx`) to that SDD vpath device. You can narrow this regular expression to only accept vpath devices and not the underlying SCSI disk devices. Modify the regular expression to accept the name vpath and to ignore all other types of devices. This is the simplest example. Adjust the example according to your environment.

```
filter = [ "a/vpath[a-z]*/", "r./*/" ]
```

This regular expression accepts all vpath devices and rejects all other devices under `/dev`.

2. Value of *types*.

In the file, you will see that it is commented out:

```
# List of pairs of additional acceptable block device types found
# in /proc/devices with maximum (non-zero) number of partitions.
# types = [ "fd", 16 ]
```

Delete the comment marker, and replace `fd` with `vpath`. This allows LVM to add vpath to its list of internally recognized devices. The partition number should stay at 16. For example:

```
types = [ "vpath", 16 ]
```

After making these two changes, save the `lvm.conf` file.

You should be able to run **pvcreate** on vpath devices (that is, `/dev/vpatha`) and create volume groups using **vgcreate**.

Booting Linux over the SAN with SDD

This section discusses the following procedures for remote booting Linux with SDD:

- “SAN Boot instructions for RHEL 3 with IBM SDD (x86)” on page 225
- “SAN Boot instructions for RHEL 3 with IBM SDD (ppc)” on page 232
- “SAN Boot Instructions for SLES 8 with IBM SDD (x86)” on page 240
- “SAN Boot Instructions for SLES 9 with IBM SDD (x86)” on page 249
- “SAN Boot instructions for SLES 9 with IBM (ppc)” on page 255
- “SAN Boot Instructions for SLES 9 with IBM SDD (x86) and LVM 2” on page 262
- “SAN boot instructions for RHEL 4 with IBM SDD (x86)” on page 270
- “SAN Boot instructions for RHEL 4 with IBM SDD (ppc)” on page 276

- “SAN boot instructions for RHEL 4 with IBM SDD (x86) and LVM 2” on page 282
- “Using lilo with SDD (remote boot) on x86” on page 289

SAN Boot instructions for RHEL 3 with IBM SDD (x86)

The following procedure is used to install RHEL 3 x86 with fibre-channel connect storage. This procedure assumes that no installation is present to work from and when completed, the boot and swap devices will be running on IBM SDD vpath devices.

Prerequisite steps

1. Ensure that the following conditions exist before continuing with this procedure:
 - a. The installation target MUST be single-pathed before installing RHEL 3.
 - b. Have a copy of RHEL 3 x86 either network-accessible or on CD-ROM.
 - c. Be familiar with the RHEL 3 installation. This includes understanding which packages will be installed.
 - d. Be familiar with how to set up a SAN network or direct-attached SAN storage devices so that the host system can access LUNs from those storage systems. (This procedure was performed on an ESS Model 800).
 - e. Be familiar with creating LUNs on the ESS Model 800 so that the host can access the ESS Model 800 devices.
 - f. Although SDD functions correctly in single-path environments, it is recommended that there be redundant physical paths to the devices from the host after installation of RHEL 3.
 - g. Optionally, have an understanding of how the Linux kernel boot process functions and what processes and procedures that are used to boot a Linux distribution for a local storage device.
 - h. Ensure that there will be network access to the system.
2. Configure QLogic Devices

Note: For ease of installation and to avoid issues with internal SCSI or IDE controllers, it is recommended that all internal disk drive controllers be disabled. This procedure assumes that this has been done.

- Verify that the QLogic SAN HBA devices that are configured for the host have been setup to have their BOOT BIOS enabled. This permits discovery and use of SAN disk devices during this procedure. While in the QLogic Utility, configure the ESS Model 800 device from which the system will boot. If the utility cannot see the correct device, check the SAN and ESS Model 800 configurations before continuing.
3. Configure Boot/Root/SWAP devices.
 - The boot device that will be used for installation and booting should be at least 4 GB in size. This is the minimum size for installing a base package set from the installation media to the boot devices.
 - It is also recommended that the swap device be at least the size of physical memory that is configured in the host. For simplicity these instructions assume that the boot, root, and swap devices are all located on the same device. However, this is not a requirement for the installation.
 4. Installation Media

The installation media; that is, the source for installation, can be CD-ROM, NFS, HTTP, FTP, and so forth. For this installation, an NFS-exported set of CD-ROMs was used. You can use any of the installation sources that are listed.

5. Install

- From the BIOS Menus select the installation source to boot from. Verify that the QLogic XXXXXXXX SAN HBA module is loaded and that the SAN devices that will be used for installation have been detected successfully.
- NOTE: Because of the way Linux discovers SAN devices, and if SAN devices have already been configured for multiple path access, Linux will discover the same physical device multiple times, once for each logical path to the device. Note which device will be used for the installation before proceeding, that is, /dev/sda.
- Select the desired options until arriving at the Installation Settings. Here, modifications of the partitioning settings are required for this installation. This is to make sure that the device noted in the previous step will be used for the root/boot installation target.
- NOTE: The details of installation and partitioning are not written up here. See the installation procedures to determine which packages are needed for the type of system being installed.

6. Rebooting

- On reboot, modify the BIOS to boot from hard disk. The system should now boot to the newly installed OS.
- Verify that the system is booted from the correct disk and vpaths.
- At this point the installed boot device can be set as the default boot device for the system. This step is not required, but is suggested because it enables unattended reboots after this procedure is complete.

7. Upgrading the SDD driver.

At the end of this document are instructions on how to upgrade the SDD driver.

Verifying the SAN boot configuration

The system should be booted in single-path mode from the SAN. Verify this by looking at the mounted devices and where the root is mounted, using the **df** command. Also verify that the swap, using **swapon -s** and other configured partitions are correctly mounted. This completes the single-path boot from SAN. To modify the boot/root and other devices for booting using the SDD driver continue with the following steps.

Note: The following instructions are examples and the values used herein might be different on your systems. In some cases, there will not be a /boot mount but rather the boot files will reside in the directory /boot on the root "/" mounted file system. It is recommended but not required that vpatha be used as the vpath boot device.

1. Install the IBM SDD driver.

Download and install the IBM SDD driver for the kernel version being used. SDD is packaged in an RPM format and can be installed using the **rpm** command. See "Installing SDD" on page 205 for more information.

2. Gather SDD data in preparation for configuring /etc/fstab, menu.lst and /boot/initrd.

```
sdd start
```

The `/etc/vpath.conf` file has now been created. You must ensure that `vpatha` is the root device. Use the `cfgvpath query` command to obtain the LUN ID of the root's physical device. (In this procedure, `sda` is the root device).

The `cfgvpath query` command produces output similar to the following example. Note that some data from the following output has been modified for ease of reading.

```
cfgvpath query
/dev/sda (8, 0) host=0 ch=0 id=0 lun=0 vid=IBM pid=2105800 serial=12020870
lun_id=12020870
/dev/sdb (8, 16) host=0 ch=0 id=0 lun=1 vid=IBM pid=2105800 serial=12120870
lun_id=12120870
/dev/sdc (8, 32) host=0 ch=0 id=0 lun=2 vid=IBM pid=2105800 serial=12220870
lun_id=12220870
/dev/sdd (8, 48) host=0 ch=0 id=0 lun=3 vid=IBM pid=2105800 serial=12320870
lun_id=12320870
```

The `lun_id` for `/dev/sda` is 12020870. Edit the `/etc/vpath.conf` file using the `lun_id` for `vpatha`. Remove all other entries from this file (they will be automatically added later by SDD).

3. Modify the `/etc/fstab` file

There is a one-to-one correlation between `sd` and `vpath` minor devices; that is, `sda1` and `vpatha1`. Major devices, however, might not necessarily correlate. For example, `sdb1` could be `vpathd1`.

Because `/boot` was installed on `/dev/sda1` and `vpatha` corresponds to `sda` in the `/etc/vpath.conf` file, `/dev/vpatha1` will be the mount device for `/boot`.

Example:

Change from:

```
LABEL=/1          /          ext3    defaults    1 1
LABEL=/boot1     /boot      ext3    defaults    1 2
none             /dev/pts   devpts  gid=5,mode=620 0 0
none             /proc      proc    defaults    0 0
none             /dev/shm   tmpfs   defaults    0 0
/dev/sda2        swap       swap    defaults    0 0
```

To:

```
/dev/vpatha3     /          ext3    defaults    1 1
/dev/vpatha1     /boot      ext3    defaults    1 2
none             /dev/pts   devpts  gid=5,mode=620 0 0
none             /proc      proc    defaults    0 0
none             /dev/shm   tmpfs   defaults    0 0
/dev/vpatha2     swap       swap    defaults    0 0
```

4. Prepare the `initrd` file.

The `[initrd file]` refers to the current `initrd` in `/boot`. The correct `initrd` can be determined by the following method:

```
ls -lA /boot | grep initrd | grep $(uname -r)
cd /boot
cp [initrd file] to initrd.vp.gz
gunzip initrd.vp.gz
mkdir /boot/mnt
```

5. For `ext2` file system `initrds`, you might need to resize the `initrd` file system.

```
dd if=/dev/zero of=initrd.vp seek=33554432 count=1 bs=1
losetup /dev/loop0 initrd.vp
e2fsck -f /dev/loop0
resize2fs -f /dev/loop0
losetup -d /dev/loop0
```

Note: Adding the `ramdisk_size=` option to the kernel entry in the boot loader file is required after increasing the size of the `initrd` file. For resizing the `initrd` to 33554432 add the following to the `/boot/grub/menu.lst` file,
`ramdisk_size=34000`

Modify the `/boot/grub/menu.lst` file. Add an entry for the SDD boot using `initrd.vp`.

```
title Red Hat Enterprise Linux AS (2.4.21-32.0.1.ELsmp) with vpath/SDD
root (hd0,0)
kernel /vmlinuz-2.4.21-32.0.1.ELsmp ro root=/dev/vpatha3 ramdisk_size=34000
initrd /initrd.vp
```

6. Change directory to `/boot` and un-archive the `initrd` image to `/boot/mnt`.
Mount the `initrd` file system.

```
mount -o loop -t ext2 initrd.vp /boot/mnt
```

7. Make additional directories in `/boot/mnt`.

```
cd /boot/mnt
mkdir mnt
mkdir lib/tls
mkdir -p opt/IBMsdd/bin
chmod -R 640 opt/IBMsdd
```

8. Copy the following `/etc` files to `/boot/mnt/etc`.

```
cp /etc/fstab /boot/mnt/etc/
cp /etc/vpath.conf /boot/mnt/etc/
cp /etc/group /boot/mnt/etc/
cp /etc/passwd /boot/mnt/etc/
cp /etc/nsswitch.conf /boot/mnt/etc/
```

9. Modify the `/boot/mnt/etc/nsswitch.conf` file if it is not as stated below.

- a. Change:

```
passwd: compat
```

To:

```
passwd: files
```

- b. Change:

```
group: compat
```

To:

```
group: files
```

10. Copy required library files for `cfgvpath`. Use the `ldd` command to determine the library files and locations.

Example:

```
ldd /opt/IBMsdd/bin/cfgvpath | awk '{print $(NF-1)}' | grep lib
```

The **ldd** command returns something similar to the following:

```
/lib/tls/libc.so.6 XXXX  
/lib/ld-linux.so.2 XXXX
```

These files must be copied to the `/boot/mnt/lib/tls/` and `/boot/mnt/lib/` directories respectively.

11. Copy the correct `sdd-mod` to the `initrd` file system.

Use the **uname -r** command to determine the correct `sdd-mod` and create a soft link.

Example:

The command will return something similar to 2.4.21-32.0.1.ELsmp

```
cp /opt/IBMsdd/sdd-mod.o-`uname -r` /boot/mnt/lib/  
cd lib  
ln -s sdd-mod.o sdd-mod.o-`uname -r` sdd-mod.o  
cd ../
```

12. Copy the following binaries and change permissions to 755:

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/  
cp /bin/awk /boot/mnt/bin/  
cp /bin/cat /boot/mnt/bin/  
cp /bin/tar /boot/mnt/bin/  
cp /bin/grep /boot/mnt/bin/  
cp /bin/chmod /boot/mnt/bin/  
cp /bin/chown /boot/mnt/bin/  
cp /bin/mknod /boot/mnt/bin/  
cp /bin/mount /boot/mnt/bin/  
cp /bin/ls /boot/mnt/bin/  
cp /bin/umount /boot/mnt/bin/  
cp /bin/cp /boot/mnt/bin/  
cp /bin/ash /boot/mnt/bin  
cp /bin/rm /boot/mnt/bin  
cp /bin/sh /boot/mnt/bin  
cp /bin/ps /boot/mnt/bin  
cp /bin/sed /boot/mnt/bin  
cp /bin/date /boot/mnt/bin  
cp /usr/bin/cut /boot/mnt/bin
```

13. Copy required library files for each binary in step 14. Use the **ldd** command to determine the library files and locations.

Note: Many binaries use the same libraries so there might be duplications of copying. Also, copy the following libraries:

```
cp /lib/libnss_files.so.2 /boot/mnt/lib  
cp /lib/libproc.so.2.0.17 /boot/mnt/lib
```

14. Modify the `/boot/mnt/linuxrc` file. Append the following lines to then end of the `linuxrc` file.

For some storage systems with Linux 2.4 kernels and addition option must be appended to the line where the `scsi_mod` module is loaded.

Change:

```
insmod /lib/scsi_mod.o
```

To:

```
insmod scsi_mod.o max_scsi_luns=256
```

The following is the original linuxrc script in the initrd file system:

```
#!/bin/nash
echo "Loading scsi_mod.o module"
insmod /lib/scsi_mod.o
echo "Loading sd_mod.o module"
insmod /lib/sd_mod.o
echo "Loading qla2300.o module"
insmod /lib/qla2300.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
echo Creating block devices
mkdevices /dev
echo Creating root device
mkrootdev /dev/root
echo 0x0100 > /proc/sys/kernel/real-root-dev
echo Mounting root filesystem
mount -o defaults --ro -t ext3 /dev/root /sysroot
pivot_root /sysroot /sysroot/initrd
umount /initrd/proc
```

The following is the modified linuxrc script in the initrd file system:

```
#!/bin/nash
echo "Loading scsi_mod.o module"
insmod /lib/scsi_mod.o max_scsi_luns=256
echo "Loading sd_mod.o module"
insmod /lib/sd_mod.o
echo "Loading qla2300.o module"
insmod /lib/qla2300.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
echo Creating block devices
mkdevices /dev
echo Loading SDD module
insmod /lib/sdd-mod.o
echo Running cfgvpath
/opt/IBMsdd/bin/cfgvpath
echo Creating block devices
mkdevices /dev
echo Copying over device files
mount -o rw -t ext3 /dev/vpatha3 /sysroot
mkdevices /sysroot/dev
umount /sysroot
#echo Creating root device
#mkrootdev /dev/root
echo 0x0100 > /proc/sys/kernel/real-root-dev
echo Mounting root filesystem
mount -o defaults --ro -t ext3 /dev/vpatha3 /sysroot
pivot_root /sysroot /sysroot/initrd
umount /initrd/procAppend Delete Edit
```

15. Prepare the mounted initrd file system for rebooting:

```
cd /boot
umount /boot/mnt
gzip initrd.vp
mv initrd.vp.gz initrd.vp
cd /
shutdown -r now
```

16. Once booted, verify that vpath devices are being used. Add all other paths and reboot again. The following commands can be used to verify the use of vpath devices.

```
mount
swapon -s
lsvpcfg
datapath query device
```

Upgrading the SDD

Use the following procedure to upgrade SDD:

1. Place the SDD rpm update package on the system
2. Restart the system in single-path mode (no SDD)
3. You might come to a **Repair filesystem** prompt. If so, then enter the root password. At the prompt, remount the file system as read/write.

```
mount -n -o remount,rw /
```

4. Remove the previous SDD driver.

```
rpm -e IBMsdd
```

The /etc/vpath.conf file will be saved to vpath.conf.rpmsave.

5. Install the new SDD driver.

```
rpm -ivh IBMsdd-x.x.x.x-y.i686.rhe13.rpm
cd /boot
mv initrd.vp initrd.vp.gz
gunzip initrd.vp.gz
mount -o loop -t ext2 initrd.vp mnt
cp /opt/IBMsdd/sdd-mod.ko-`uname -r` /boot/mnt/lib/
```

6. Verify that the soft link sdd-mod.ko in /boot/mnt/lib points to the current sdd module.
7. Copy the new **cfgvpath** command and use the **ldd** command to verify that the correct libraries are installed for /boot/mnt/opt/IBMsdd/bin/cfgvpath.

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
```

8. Prepare the initrd.vp for booting.

```
umount mnt
gzip initrd.vp
mv initrd.vp.gz initrd.vp
cd /
```

9. Reboot the system.

SAN Boot instructions for RHEL 3 with IBM SDD (ppc)

The following procedure is used to install RHEL 3 ppc onto a IBM System p JS20 LPAR with fibre-channel connected storage. It is assumed that no installation is present to work from and when completed, the boot and swap devices will be running on IBM SDD vpath devices.

Prerequisite steps

1. Ensure that the following conditions exist before continuing with this procedure:
 - Have a copy of RHEL 3 either network accessible or on CD-ROM.
 - Be familiar with the Red Hat installation. This includes understanding which packages will be installed and how to select required options through the installation.
 - Be familiar with how to connect to and operate IBM BladeCenter control or IBM System p LPAR.
 - Be familiar with how to setup an LPAR and select the required resources to create a configured LPAR with processors, memory, and SAN HBAs. For network installs, a network port is required, and for CD-ROM installs a CD-ROM is required.
 - Be familiar with how to setup a SAN network or direct-attached SAN storage devices so that the configured LPAR can access LUNs from the storage unit.
 - Be familiar with creating LUNs on the storage unit so that the LPAR can access the storage devices. Although SDD functions correctly in single-path environments, there should be redundant physical paths to the devices from the host (after installation).
 - Optionally, have an understanding of how the Linux kernel boot process functions and what processes and procedures that are used to boot a Linux distribution for a local storage device.
2. Configure Fibre Channel Adapters
 - Verify that the SAN HBA devices that are configured for the system have been setup to have their BOOT BIOS enabled. This permits discovery and use of SAN disk devices during this procedure.
3. Configure root/boot/swap devices
 - The physical boot device that will be used for installation and booting should be at least 4 GB in size. This is the minimum size for installing all packages from the installation media to the boot devices. It is also recommended that the swap device be at least the size of physical memory that is configured in the LPAR. For simplicity these instructions assume that the root/boot/swap devices are all located on the same device; however this is not a requirement for the installation. Also, It is not required that a /boot mount exists. In some cases, there will not be a /boot mount but rather the boot files will reside in the directory /boot on the root "/" mount.
4. Installation Media
 - The installation media; that is, the source for installation, can be CD-ROM, NFS, HTTP, FTP, and so forth. For this installation, an NFS-exported set of CD-ROMs was used. You can use any of the installation sources that are listed.
5. Upgrading the SDD driver and/or OS
 - At the end of this document are instructions on how to upgrade the SDD driver.

- Each time the OS is updated or a new initrd is created, these procedures must be performed for the new OS and initrd.

Use this procedure to install RHEL 3:

1. From the SMS menu, select the installation source and boot from the media.
2. Verify that the Fibre HBA module is loaded and that the SAN devices that will be used for installation have been detected successfully.

Note: Because of the way Linux discovers SAN devices, and if SAN devices have already been configured for multiple path access, Linux will discover the same physical device multiple times, once for each logical path to the device. Take note which device will be used for the installation before proceeding, that is, `/dev/sdb`. Also note which of the fibre HBA devices is used to discover this device as it will be needed in a later step.

3. Select the desired options until arriving at the Installation Settings step of the `yast` install.

Here, modification of the partitioning settings is required for this installation. This is to make sure that the device noted in the previous step will be used for the root/boot installation target.

- a. Select partitioning, and go to the custom partition setup.
- b. Select the device and Custom partitioning for experts.
- c. Make sure that there is a PReP boot partition on the root/boot device and that it is the first partition.
- d. Continue to partition devices as required to complete this configuration.

The details of installation and partitioning are not written up here. See the installation procedures to determine which packages are needed for the type of system being installed.

4. Finish the installation. If an error occurs while attempting to create the `yaboot` boot loader stating that the device type of `fc` is unknown. Select **OK** and select **No** when asked to retry.
5. Reboot the SMS menu.

This time the boot device which has been setup over the previous steps is now ready to be booted.

6. Select to boot from a Hard Drive/SAN and select the Fibre HBA device adapter associated with the SAN disk device which the installation was completed.

The installation boot device should now be listed in the bootable devices discovered in the SAN on the selected Fibre HBA.

7. Select the appropriate device and boot.

SAN boot configuration

The system should now be booted in single-path mode from the SAN. Verify this by looking at the mounted devices and where the root is mounted, using the `df` command. Also verify that the swap, using `swapon -s` and other configured partitions are correctly mounted. This completes the single-path boot from SAN. To modify the boot/root and other devices for booting using the SDD driver continue with the following steps.

Note: The following instructions are examples and the values used herein might be different on your systems. In some cases, there will not be a `/boot` mount but rather the boot files will reside in the directory `/boot` on the root `"/"` mounted file system. It is recommended but not required that `vpatha` be used as the `vpath` boot device.

1. Install the IBM SDD driver.

Download and install the IBM SDD driver for the kernel version being used. SDD is packaged in an RPM format and can be installed using the **rpm** command. See “Installing SDD” on page 205 for more information.

2. Gather SDD data in preparation for configuring `/etc/vpath.conf`, `/etc/fstab/`, `etc/yaboot.conf` and `/boot/initrd`.

```
sdd start
```

The `/etc/vpath.conf` file has now been created. You must ensure that `vpatha` is the root device. Use the **cfgvpath query** command to obtain the LUN ID of the root's physical device. (In this procedure, `sda` is the root device).

The **cfgvpath query** command produces output similar to the following: Note that some data from the following output has been modified for ease of reading.

```
cfgvpath query
```

```
/dev/sda ( 8, 0) host=0 ch=0 id=0 lun=0 vid=IBM pid=2105800
serial=13320870 lun_id=13320870 ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdb not configured: Either in /etc/fstab, or mounted or is a raw device
/dev/sdb ( 8, 16) host=0 ch=0 id=0 lun=1 vid=IBM pid=2105800
serial=13E20870 lun_id=13E20870 ctlr_flag=0 ctlr_nbr=0 df_ctlr=0 X
/dev/sdc ( 8, 32) host=0 ch=0 id=0 lun=2 vid=IBM pid=2105800
serial=12E20870 lun_id=12E20870 ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdd not configured: Either in /etc/fstab, or mounted or is a raw device
/dev/sdd ( 8, 48) host=0 ch=0 id=0 lun=3 vid=IBM pid=2105800
serial=13F20870 lun_id=13F20870 ctlr_flag=0 ctlr_nbr=0 df_ctlr=0 X
```

The `lun_id` for `/dev/sdb` is `13E2087`. Edit the `/etc/vpath.conf` file using the `lun_id` for `vpatha` (`vpatha 13E2087`). Remove all other entries from this file (they will be automatically added later by SDD) .

Contents of `/etc/vpath.conf` :

```
vpatha 13E280
```

3. Extracting and mounting the `initrd`

The following unzips and extracts the `initrd` image so that it can be modified to include the required elements to enable a `vpath` boot image.

```
cd /boot
```

Locate the `initrd` image used for booting. This will be the image that `/etc/yaboot.conf` is pointing to. Note that the file pointed to might be symbolic link to another file. Copy the file to a temporary filename with a `.gz` extension. For example, if the file name is `initrd-2.4.21-32.0.1.EL.img`, the correct [initrd file] can be determined by the following method:

```
cd /boot
ls -lA /boot | grep initrd | grep $(uname -r)
cp [initrd file] to initrd.vp.gz
gunzip initrd.vp.gz
mkdir /boot/mnt
```

Create a temporary directory where the image will be manipulated, for example, `/boot/mnt`, This is referred to as the *image temporary directory* throughout the rest of this documentation.

```
mkdir /boot/mnt
```

For ext2 file system initrd's, you might be required to resize the initrd file system (recommended).

```
dd if=/dev/zero of=initrd.vp seek=33554432 count=1 bs=1
e2fsck -f /boot/initrd.vp
```

Output similar to the following will be displayed. Answer y when prompted:

```
e2fsck 1.32 (09-Nov-2002)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
/lost+found not found. Create <y>? y

Pass 4: Checking reference counts
Pass 5: Checking group summary information

initrd.vp: ***** FILE SYSTEM WAS MODIFIED *****
initrd.vp: 36/2000 files (0.0% non-contiguous), 2863/8000 blocks
[root@elml7a212 boot]#

resize2fs -f /boot/initrd.vp
```

Note: Adding the `ramdisk_size=` option to the kernel entry in the boot loader file might be required after increasing the size of the initrd file. For resizing the initrd to 33554432 add the following to the `/boot/grub/menu.lst` file in the “append section”, `ramdisk_size=34000`. An example of this entry is provided later in this topic.

Mount the initrd file system.

```
mount -o loop -t ext2 initrd.vp /boot/mnt
```

4. Modifying the `/boot/initrd`

Create the following directories in the image temporary directory.

```
cd /boot/mnt
mkdir mnt
mkdir -p lib/tls
mkdir -p lib64/tls
mkdir -p opt/IBMsdd/bin
chmod -R 640 opt/IBMsdd
```

Copy the following files to the following directories relative to the image temporary directory.

```

cp /opt/IBMsdd/sdd-mod.o-`uname -r` /boot/mnt/lib/
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
cp /bin/awk /boot/mnt/bin/
cp /bin/cat /boot/mnt/bin/
cp /bin/tar /boot/mnt/bin/
cp /bin/grep /boot/mnt/bin/
cp /bin/chmod /boot/mnt/bin/
cp /bin/chown /boot/mnt/bin/
cp /bin/mknod /boot/mnt/bin/
cp /bin/mount /boot/mnt/bin/
cp /bin/ls /boot/mnt/bin/
cp /bin/umount /boot/mnt/bin/
cp /bin/cp /boot/mnt/bin/
cp /bin/ash /boot/mnt/bin
cp /bin/rm /boot/mnt/bin
cp /bin/sh /boot/mnt/bin
cp /bin/ps /boot/mnt/bin
cp /bin/sed /boot/mnt/bin
cp /bin/date /boot/mnt/bin
cp /usr/bin/cut /boot/mnt/bin

```

Issue the following command from the lib directory. The linked module is the name of the module that was copied into the /boot/mnt/lib directory above.

```

cd /boot/mnt/lib
ln -s sdd-mod.o-2.4.21-32.0.1.EL sdd-mod.o

```

For each of the above binary files (except sdd-mod.o), run the **ldd** command and verify that the listed library files exist in the image temporary directory. If they do not, copy the listed library files that do not exist to the corresponding lib and lib64 directories in the image temporary directory.

5. Copy required library files for 'cfgvpath'. Use the 'ldd' command to determine the library files and locations.

Example:

```

ldd /opt/IBMsdd/bin/cfgvpath | awk '{print $(NF-1)}' | grep lib

```

The **ldd** command might return something similar to the following:

```

/lib64/tls/libc.so.6 XXXX
/lib64/ld-linux.so.2 XXXX

```

These file must be copied to the /boot/mnt/lib64/tls/ and /boot/mnt/lib64/ directories respectively.

Copy this additional library file:

```

cp /lib/libnss_files.so.2 /boot/mnt/lib

```

6. Copy /etc files to the image temporary directory /boot/mnt/etc.

```

cp /etc/vpath.conf /boot/mnt/etc
cp /etc/passwd /boot/mnt/etc
cp /etc/group /boot/mnt/etc
cp /etc/nsswitch.conf /boot/mnt/etc

```

7. Modify the /etc/fstab to use vpath devices for /root and swap.

Other devices using vpaths will also need to be changed. For the initial install it is recommended to only work with the root/boot/swap devices and comment out other sd and hd devices until completed.

Original:

```

LABEL=/1      /          ext3  defaults  1 1
none         /dev/pts  devpts gid=5,mode=620 0 0
none         /proc     proc  defaults  0 0
none         /dev/shm  tmpfs defaults  0 0
/dev/sdd3    swap      swap  defaults  0 0
/dev/sdb3    swap      swap  defaults  0 0

```

Modified:

```

/dev/vpatha2  /          ext3  defaults  1 1
none         /dev/pts  devpts gid=5,mode=620 0 0
none         /proc     proc  defaults  0 0
none         /dev/shm  tmpfs defaults  0 0
#/dev/sdd3    swap      swap  defaults  0 0
/dev/vpatha3  swap      swap  defaults  0 0

```

```
cp /etc/fstab /boot/mnt/etc
```

8. Edit the /boot/mnt/linuxrc file.

For some storage systems with Linux 2.4 kernels an additional option must be appended to the line where the `scsi_mod` module is loaded.

Change:

```
insmod /lib/scsi_mod.o
```

To:

```
insmod scsi_mod.o max_scsi_luns=256
```

Original linuxrc script in the `initrd` file system

```

#!/bin/nash
echo "Loading scsi_mod.o module"
insmod /lib/scsi_mod.o
echo "Loading sd_mod.o module"
insmod /lib/sd_mod.o
echo "Loading qla2300.o module"
insmod /lib/qla2300.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
echo Creating block devices
mkdevices /dev
echo Creating root device
mkrootdev /dev/root
echo 0x0100 > /proc/sys/kernel/real-root-dev
echo Mounting root filesystem
mount -o defaults --ro -t ext3 /dev/root /sysroot
pivot_root /sysroot /sysroot/initrd
umount /initrd/proc

```

Modified linuxrc script in the `initrd` file system

```

#!/bin/nash
echo "Loading scsi_mod.o module"
insmod /lib/scsi_mod.o max_scsi_luns=256
echo "Loading sd_mod.o module"
insmod /lib/sd_mod.o
echo "Loading qla2300.o module"
insmod /lib/qla2300.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
echo Creating block devices
mkdevices /dev
echo "Loading SDD module"
insmod /lib/sdd-mod.o
echo "Configuring vpath devices"
/opt/IBMsdd/bin/cfgvpath
echo Copying over device files
mount -o rw -t ext3 /dev/vpatha3 /sysroot
mkdevices /sysroot/dev
umount /sysroot
echo 0x0100 > /proc/sys/kernel/real-root-dev
echo Mounting root filesystem
mount -o defaults --ro -t ext3 /dev/vpatha3 /sysroot
pivot_root /sysroot /sysroot/initrd
umount /initrd/proc

```

9. Repackaging the initrd.

To repackage all of the changes that have just been made to the initrd, issue the following commands:

```

cd /boot
umount /boot/mnt
gzip initrd.vp
mv initrd.vp.gz initrd.vp

```

The `initrd-2.4.21-32.0.1.EL.img` now has the repackaged initrd image with the SDD drive and modified files required to boot from a vpath.

10. Modifying root device files and updating the boot partition.

Modify `/etc/yaboot.conf`. Add a new entry in the file and modify the entry to point at the new initrd image created in the previous step. Also modify the root device in the new entry to point to the vpath chosen from the previous steps. Remember to include the partition if required. Also make sure to modify the entry name.

Original `/etc/yaboot.conf`:

```

image=/boot/vmlinux-2.4.21-32.0.1.EL
label=2.4.21-32.0.1.E
read-only
initrd=/boot/initrd-2.4.21-32.0.1.EL.img
append="console=hvc0 root=/LABEL/"

```

Modified `/etc/yaboot.conf`:

```

image=/boot/vmlinuz-2.4.21-32.0.1.EL
label=2.4.21-32.0.1.E
read-only
initrd=/boot/initrd-2.4.21-32.0.1.EL.img
append="console=hvc0 root=/LABEL=/"
image=/boot/vmlinuz-2.4.21-32.0.1.EL
label=2.4.21-32.0.1.E_SDD
read-only
initrd=/boot/initrd.vp
append="console=hvc0 root=/dev/vpatha3 ramdisk_size=34000"

```

Write the yaboot configuration to the boot partition:

```
ybin -b /dev/sdb1
```

11. Restart the system
 - a. Reboot the system.
 - b. From the SMS menu, select the boot devices as you did before, if the boot device is not already set up as the first boot device.
 - c. When the yaboot prompt is shown during boot, type in the given name for the new boot image.
 - d. During the OS load, ensure that the IBMsdd module is loaded after the SAN disk devices are discovered.
 - e. Ensure that there were no errors printed to the console during boot.
 - f. If there were errors, reboot the system and at the yaboot prompt, select the old image to boot from. When the system boots, review the preceding steps and make any corrections to errors, then repeat these steps, starting with step 9 (repackaging the initrd).

12. Verify System has reboot and SDD is configured correctly

Once booted, verify that vpath devices are being used. Add all other paths and reboot again. The following commands can be used to verify the use of vpath devices:

- mount
- swapon -s
- lsvpcfg
- datapath query device

At this point, the installed boot device can be set as the default boot device for the system via the SMS menu. This step is not required, but is suggested because it enables unattended reboots after this procedure is complete.

SDD upgrade procedure

Use the following procedure to upgrade SDD:

1. Place the SDD rpm update package on the system
2. Restart the system in single-path mode (no SDD)
3. You might come to a **Repair filesystem** prompt. If so, enter the root password. At the prompt, remount the file system as read/write.

```
mount -n -o remount,rw /
```

For systems that have /boot on a separate mount point, mount /boot partition using /dev/sd device.

4. Remove the previous SDD driver.

```
rpm -e IBMsdd
```

The `/etc/vpath.conf` file will be saved to `vpath.conf.rpmsave`.

5. Install the new SDD driver.

```
rpm -ivh IBMsdd-x.x.x.x-y.ppc64.rhel3.rpm
cd /boot
mv initrd.vp initrd.vp.gz
gunzip initrd.vp.gz
mount -o loop -t ext2 initrd.vp mnt
cp /opt/IBMsdd/sdd-mod.ko-`uname -r` /boot/mnt/lib/
```

6. Verify that the soft link `sdd-mod.ko` in `/boot/mnt/lib` points to the current `sdd` module.
7. Copy the new **cfgvpath** command and use the **ldd** command to verify that the correct libraries are installed for `/boot/mnt/opt/IBMsdd/bin/cfgvpath`.

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
```

8. Prepare the `initrd.vp` for booting.

```
umount mnt
gzip initrd.vp
mv initrd.vp.gz initrd.vp
cd /
```

9. Reboot the system.

SAN Boot Instructions for SLES 8 with IBM SDD (x86)

The following procedure is used to install SLES 8 SP4 i386 on an xSeries host with fibre-channel connect storage. This procedure assumes that no installation is present to work from and when completed, the boot and swap devices will be running on IBM SDD vpath devices. This procedure is intended to be used with the following revisions of software, firmware, and hardware. Variations from the documents have not been specifically tested.

Prerequisite steps

1. Ensure that the following conditions exist before continuing with this procedure:
 - a. The installation target **MUST** be single-pathed before installing SLES 8.
 - b. Have a copy of SLES 8 SP4 i386 either network-accessible or on CD-ROM.
 - c. Be familiar with the SLES 8 installation. This includes understanding which packages will be installed.
 - d. Be familiar with how to set up a SAN network or direct-attached SAN storage devices so that the host system can access LUNs from those storage systems. (This procedure was performed on an ESS Model 800).
 - e. Be familiar with creating LUNs on the ESS Model 800 so that the host can access the ESS Model 800 devices.
 - f. Although SDD functions correctly in single-path environments, it is recommended that there be redundant physical paths to the devices from the host after installation of SLES 8.
 - g. Optionally, have an understanding of how the Linux kernel boot process functions and what processes and procedures that are used to boot a Linux distribution for a local storage device.

- h. Ensure that there will be network access to the system.
2. Configure QLogic Devices
 - For ease of installation and to avoid issues with internal SCSI or IDE controllers, it is recommended that all internal disk drive controllers be disabled. This procedure assumes that this has been done.
 - Verify that the QLogic SAN HBA devices that are configured for the host have been setup to have their BOOT BIOS enabled. This permits discovery and use of SAN disk devices during this procedure. While in the QLogic Utility, configure the ESS Model 800 device from which the system will boot. If the utility cannot see the correct device, check the SAN and ESS Model 800 configurations before continuing.
3. Configure Boot/Root/SWAP devices

The boot device that will be used for installation and booting should be at least 4 GB in size. This is the minimum size for installing a base package set from the installation media to the boot devices.

The swap device be at least the size of physical memory that is configured in the host. For simplicity, these instructions assume that the boot, root, and swap devices are all located on the same device; however, this is not a requirement for the installation.
4. Installation Media

The installation media; that is, the source for installation, can be CD-ROM, NFS, HTTP, FTP, and so forth. For this installation an NFS-exported set of CD-ROMs was used. You can use any of the installation sources that are listed.
5. Install
 - From the BIOS Menus select the installation source to boot from. Verify that the QLogic XXXXXXXX SAN HBA module is loaded and that the SAN devices that will be used for installation have been detected successfully.
 - Because of the way Linux discovers SAN devices, and if SAN devices have already been configured for multiple path access, Linux will discover the same physical device multiple times, once for each logical path to the device. Note which device will be used for the installation before proceeding, for example, /dev/sda.
 - Select the desired options until arriving at the Installation Settings. Here, modifications of the partitioning settings are required for this installation. This is to make sure that the device noted in the previous step will be used for the root/boot installation target.
 - The details of installation and partitioning are not included here. See the installation procedures to determine which packages are needed for the type of system being installed.
6. Rebooting
 - a. On reboot, modify the BIOS to boot from hard disk. The system should now boot to the newly installed OS.
 - b. Verify that the system is booted from the correct disk and vpaths.
 - c. At this point the installed boot device can be set as the default boot device for the system. This step is not required, but is suggested because it enables unattended reboots after this procedure is complete.
7. Upgrading the SDD driver

At the end of this document, are instructions on how to upgrade the SDD driver.

SAN boot configuration

The system should be booted in single-path mode from the SAN. Verify this by looking at the mounted devices and where the root is mounted, using the **df** command. Also verify that the swap, using **swapon -s** and other configured partitions are correctly mounted. This completes the single-path boot from SAN. To modify the boot/root and other devices for booting using the SDD driver continue with the following steps.

Note: The following instructions are examples and the values used herein might be different on your systems. Also, any time a kernel is upgraded a new initrd for that kernel must be built following the instructions in the procedure.

1. Install the IBM SDD driver.

Download and install the IBM SDD driver for the kernel version being used. SDD is packaged in an RPM format and can be installed using the **rpm** command. See “Installing SDD” on page 205 for more information.

2. Gather SDD data in preparation for configuring `/etc/fstab`, `menu.lst` and `/boot/initrd`.

```
sdd start
```

The `/etc/vpath.conf` file has now been created. You must ensure that `vpatha` is the root device. Use the **cfgvpath query** command to get the LUN ID of the root's physical device. (In this procedure, `sda` is the root device).

The **cfgvpath query** command will produce output similar to the following: Note that some data from the following output has been modified for ease of reading.

```
cfgvpath query
/dev/sda (8, 0) host=0 ch=0 id=0 lun=0 vid=IBM pid=2105800 serial=12020870
lun_id=12020870
/dev/sdb (8, 16) host=0 ch=0 id=0 lun=1 vid=IBM pid=2105800 serial=12120870
lun_id=12120870
/dev/sdc (8, 32) host=0 ch=0 id=0 lun=2 vid=IBM pid=2105800 serial=12220870
lun_id=12220870
/dev/sdd (8, 48) host=0 ch=0 id=0 lun=3 vid=IBM pid=2105800 serial=12320870
lun_id=12320870
```

The `lun_id` for `/dev/sda` is 12020870. Edit the `/etc/vpath.conf` file using the `lun_id` for `vpatha`. Remove all other entries from this file (they will be automatically added later by SDD).

3. Modify the `/etc/fstab` and the `/boot/grub/menu.lst`

There is a one-to-one correlation between `sd` and `vpath` minor devices, such as, `sda1` and `vpatha1`.

Major devices, however, might not necessarily correlate; for example, `sdb1` could be `vpathd1`.

Because `/boot` was installed on `/dev/sda1` and `vpatha` corresponds to `sda` in the `/etc/vpath.conf` file, `/dev/vpatha1` will be the mount device for `/boot`.

Example:

Change from:

```

/dev/sda3      /          ext3  defaults  1 1
/dev/sda1      /boot      ext3  defaults  1 2
none          /dev/pts   devpts gid=5,mode=620 0 0
none          /proc      proc  defaults  0 0
none          /dev/shm   tmpfs  defaults  0 0
/dev/sda2      swap       swap  defaults  0 0

```

To:

```

/dev/vpatha3   /          ext3  defaults  1 1
/dev/vpatha1   /boot      ext3  defaults  1 2
none          /dev/pts   devpts gid=5,mode=620 0 0
none          /proc      proc  defaults  0 0
none          /dev/shm   tmpfs  defaults  0 0
/dev/vpatha2   swap       swap  defaults  0 0

```

Modify the `/boot/grub/menu.lst` file. Add an entry for the SDD boot using `initrd.vp`

```

title linux-smp
kernel (hd0,0)/vmlinuz-2.4.21-295-smp root=/dev/sda3
initrd (hd0,0)/initrd-2.4.21-295-smp
title linux-smp-SDD
kernel (hd0,0)/vmlinuz-2.4.21-295-smp root=/dev/vpatha3 ramdisk_size=34000
initrd (hd0,0)/initrd.vp

```

4. Prepare the `initrd` file.

The `[initrd file]` refers to the current `initrd` in `/boot`. The correct `initrd` can be determined by the following method:

```

ls -lA /boot | grep initrd | grep $(uname -r)
cd /boot
cp [initrd file] to initrd.vp.gz
gunzip initrd.vp.gz
mkdir /boot/mnt

```

5. For `ext2` file system `initrds`, you might need to resize the `initrd` file system. For `Sles8u5`, this step might not be required.

```

dd if=/dev/zero of=initrd.vp seek=33554432 count=1 bs=1
losetup /dev/loop0 initrd.vp
e2fsck -f /dev/loop0
resize2fs -f /dev/loop0
losetup -d /dev/loop0

```

Note: Adding the `ramdisk_size=` option to the kernel entry in the boot loader file might be required after increasing the size of the `initrd` file. For resizing the `initrd` to 33554432 add the following to the `/boot/grub/menu.lst` file, `ramdisk_size=34000` (see the previous step for modifying the `/boot/grub/menu.lst`).

6. Change directory to `/boot` and un-archive the `initrd` image to `/boot/mnt`. Mount the `initrd` file system.

```

mount -o loop -t ext2 initrd.vp /boot/mnt

```

7. Make additional directories in `/boot/mnt` if they do not exist.

```
cd /boot/mnt
mkdir mnt
mkdir -p opt/IBMsdd/bin
chmod -R 640 opt/IBMsdd
mkdir -p lib/tls (For some Sles8 version create lib/i686)
mkdir sysroot
mkdir proc
```

8. Copy the following /etc files to /boot/mnt/etc.

```
cp /etc/fstab /boot/mnt/etc/
cp /etc/vpath.conf /boot/mnt/etc/
cp /etc/group /boot/mnt/etc/
cp /etc/passwd /boot/mnt/etc/
cp /etc/nsswitch.conf /boot/mnt/etc/
```

9. Modify the /boot/mnt/etc/nsswitch.conf file.

- a. Change:

```
passwd: compat
```

To:

```
passwd: files
```

- b. Change:

```
group: compat
```

To:

```
group: files
```

10. Copy cfgvpath to the initrd image.

Copy /opt/IBMsdd/bin/cfgvpath to /boot/mnt/opt/IBMsdd/bin/ and change permissions of cfgvpath to 755.

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
```

11. Copy required library files for cfgvpath .

Use the **ldd** command to determine the library files and locations.

Example:

```
ldd /opt/IBMsdd/bin/cfgvpath | awk {print $(NF-1)} | grep lib
```

The output of the **ldd** command might be similar to the following:

```
/lib/i686/libc.so.6 XXXX
/lib/ld-linux.so.2 XXXX
```

These file must be copied to the /boot/mnt/lib/i686/ and /boot/mnt/lib/ directories respectively.

12. Copy the correct sdd-mod.o file to the initrd file system. Use the **uname -r** command to determine the correct sdd-mod.o file and create a soft link.

Example:

The **uname -r** command will return something similar to 2.6.5-7.191-smp.

```
cp /opt/IBMsdd/sdd-mod.o-`uname -r` /boot/mnt/lib/  
cd lib  
ln -s sdd-mod.o sdd-mod.o-`uname -r` sdd-mod.o  
cd ../
```

13. Copy the following binaries to the initrd bin directory.

```
cp /bin/awk /boot/mnt/bin/  
cp /bin/chmod /boot/mnt/bin/  
cp /bin/chown /boot/mnt/bin/  
cp /bin/cp /boot/mnt/bin/  
cp /bin/date /boot/mnt/bin/  
cp /bin/grep /boot/mnt/bin/  
cp /bin/ls /boot/mnt/bin/  
cp /bin/mknod /boot/mnt/bin/  
cp /bin/mount /boot/mnt/bin/  
cp /bin/ps /boot/mnt/bin/  
cp /bin/rm /boot/mnt/bin/  
cp /bin/sed /boot/mnt/bin/  
cp /bin/sh /boot/mnt/bin/  
cp /bin/tar /boot/mnt/bin/  
cp /bin/umount /boot/mnt/bin/  
cp /bin/cat /boot/mnt/bin/  
  
cp /usr/bin/cut /boot/mnt/bin/  
cp /usr/bin/expr /boot/mnt/bin/
```

14. Copy the required library files for each binary in the previous step.
Use the **ldd** command to determine the library files and locations.

Note: Many binaries use the same libraries so there might be duplications of copying.

Also, copy the following library.

```
cp /lib/libnss_files.so.2 /boot/mnt/lib
```

15. Copy /dev/sd devices to the initrd /dev directory using the **tar** command.

```
cd /boot/mnt  
tar cps /dev/sd[a-z] /dev/sd[a-z][a-z] | tar xps
```

16. Modify the /boot/mnt/linuxrc file.

Add the following lines just after the last kernel module is loaded.

For 2.4 kernels an addition option must be appended to the line where the `scsi_mod` module is loaded for storage systems such as the DS6000 and the DS8000.

Change:

```
insmod /lib/scsi_mod.o
```

To:

```
insmod scsi_mod.o max_scsi_luns=256
```

Add the following lines to the linuxrc file after the last driver has been loaded.

```

echo "Mounting proc"
mount -n -tproc none /proc

echo "Loading SDD module"
insmod /lib/sdd-mod.o

echo "Running SDD configuration"
/opt/IBMsdd/bin/cfgvpath

```

Ensure an updated copy of vpath.conf and the vpath device files are copied to the root file system during boot by using the following syntax to mount the root file system.

```
mount -o rw -t [fstype] [device] /mnt
```

Add the following lines just after the modules load entries. The values used for the [fstype] and [device] here are only examples. Use the correct values for the system that is being configured.

```

echo "Copying over device files"
mount -o rw -t ext3 /dev/vpatha3 /sysroot
(tar cps /dev/IBMsdd /dev/vpath*) | (cd /sysroot && tar xps)
cp /etc/vpath.conf /sysroot/etc/
umount /sysroot

```

You must ensure that the correct major and minor number of the /root vpath device is written to /proc/sys/kernel/real-root-dev. To do this, add the following line to the linuxrc file.

```

echo "Setting correct root device"
for name in `cat /proc/cmdline`; do
#Look for "root="
echo $name | grep -q ^root
if [ $? -eq 0 ]; then
# chop off the "root="
dev_name=`expr "$name" : .*=(.*) `
echo "Found root = $dev_name"
#chop off the "dev"
dev_name=`expr "$dev_name" : /dev/(.*) `
#find the major/minor in /proc/partitions
parts=`grep $dev_name /proc/partitions`
dev_major=`echo $parts | cut -d -f1`
dev_minor=`echo $parts | cut -d -f2`
dev_num=`expr $dev_major \* 256 + $dev_minor`
echo $dev_num > /proc/sys/kernel/real-root-dev
continue
fi
done

```

Original linuxrc script in the initrd file system:

```

#!/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/2.4.21-295-smp/kernel/drivers/scsi/scsi_mod.o $extra_scsi_params

echo "Loading kernel/drivers/scsi/sd_mod.o"
insmod /lib/modules/2.4.21-295-smp/kernel/drivers/scsi/sd_mod.o

echo "Loading kernel/fs/jbd/jbd.o"
insmod /lib/modules/2.4.21-295-smp/kernel/fs/jbd/jbd.o

echo "Loading kernel/fs/ext3/ext3.o"
insmod /lib/modules/2.4.21-295-smp/kernel/fs/ext3/ext3.o

echo "Loading kernel/drivers/scsi/qla2300.o"
insmod /lib/modules/2.4.21-295-smp/kernel/drivers/scsi/qla2300.o

echo "Loading kernel/drivers/scsi/qla2300_conf.o"
insmod /lib/modules/2.4.21-295-smp/kernel/drivers/scsi/qla2300_conf.o

Modified linuxrc script in the initrd file system with modifications.

#!/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 max_scsi_luns=255"
insmod /lib/modules/2.4.21-295-smp/kernel/drivers/scsi/scsi_mod.o
(this line is part of the above line) $extra_scsi_params max_scsi_luns=255

echo "Loading kernel/drivers/scsi/sd_mod.o"
insmod /lib/modules/2.4.21-295-smp/kernel/drivers/scsi/sd_mod.o

echo "Loading kernel/fs/jbd/jbd.o"
insmod /lib/modules/2.4.21-295-smp/kernel/fs/jbd/jbd.o

echo "Loading kernel/fs/ext3/ext3.o"
insmod /lib/modules/2.4.21-295-smp/kernel/fs/ext3/ext3.o

echo "Loading kernel/drivers/scsi/qla2300.o"
insmod /lib/modules/2.4.21-295-smp/kernel/drivers/scsi/qla2300.o

echo "Loading kernel/drivers/scsi/qla2300_conf.o"
insmod /lib/modules/2.4.21-295-smp/kernel/drivers/scsi/qla2300_conf.o

```

```

echo "Mounting proc"
mount -n -tproc none /proc

echo "Loading SDD module"
insmod /lib/sdd-mod.o

echo "Running SDD configuration"
/opt/IBMsdd/bin/cfgvpath

echo "Copying over device files"
mount -o rw -t ext3 /dev/vpatha3 /sysroot
(tar cps /dev/IBMsdd /dev/vpath*) | (cd /sysroot && tar xps)
umount /sysroot

echo "Setting correct root device"
for name in `cat /proc/cmdline`; do
    #Look for "root="
    echo $name | grep -q ^root
    if [ $? -eq 0 ]; then
        # chop off the "root="
        dev_name=`expr "$name" : .*=\(.*\)`
        echo "Found root = $dev_name"
        #chop off the "dev"
        dev_name=`expr "$dev_name" : /dev/\(.*\)`
        #find the major/minor in /proc/partitions
        parts=`grep $dev_name /proc/partitions`
        dev_major=`echo $parts | cut -d -f1`
        dev_minor=`echo $parts | cut -d -f2`
        dev_num=`expr $dev_major \* 256 + $dev_minor`
        echo $dev_num > /proc/sys/kernel/real-root-dev
        continue
    fi
done

echo "Unmounting proc"
umount /proc

```

17. Prepare the mounted initrd file system for rebooting:

```

cd /boot
umount /boot/mnt
gzip initrd.vp
mv initrd.vp.gz initrd.vp
cd /
shutdown -r now

```

18. Once booted, verify that vpath devices are being used. Add all other paths and reboot again. The following commands can be used to verify the use of vpath devices.

```

mount
swapon -s
lsvpcfg
datapath query device

```

SDD upgrade procedure

Use the following procedure to upgrade SDD:

1. Place the SDD rpm update package on the system
2. Restart the system in single-path mode (no SDD)
3. You might come to a **Repair filesystem** prompt. If so, enter the root password. At the prompt, remount the file system as read/write.

```
mount -n -o remount,rw /
```

4. For systems that have /boot on a separate mount point, mount /boot partition using /dev/sa device.

5. Remove the previous SDD driver.

```
rpm -e IBMsdd
```

The `/etc/vpath.conf` file will be saved to `vpath.conf.rpmsave`.

6. Install the new SDD driver.

```
rpm -ivh IBMsdd-x.x.x.x-y.i686.sles8.rpm
cd /boot
mv initrd.vp initrd.vp.gz
gunzip initrd.vp.gz
mount -o loop -t ext2 initrd.vp mnt
cp /opt/IBMsdd/sdd-mod.ko-`uname -r` /boot/mnt/lib/
```

7. Verify that the soft link `sdd-mod.ko` in `/boot/mnt/lib` points to the current SDD module.
8. Copy the new **cfgvpath** command and use the **ldd** command to verify that the correct libraries are installed for `/boot/mnt/opt/IBMsdd/bin/cfgvpath`.

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
```

9. Prepare the `initrd.vp` for booting.

```
umount mnt
gzip initrd.vp
mv initrd.vp.gz initrd.vp
cd /
```

10. Reboot the system.

SAN Boot Instructions for SLES 9 with IBM SDD (x86)

The following procedure is used to install SLES 9 x86 on an xSeries host with fibre-channel connect storage. This procedure assumes that no installation is present to work from and when completed, the boot and swap devices will be running on IBM SDD vpath devices.

Prerequisite steps

1. Ensure that the following conditions exist before continuing with this procedure:
 - a. The installation target MUST be single-pathed before installing SLES 9.
 - b. Have a copy of SLES 9 SP2 i386 either network-accessible or on CD-ROM.
 - c. Be familiar with the SLES 9 installation. This includes understanding which packages will be installed.
 - d. Be familiar with how to set up a SAN network or direct-attached SAN storage devices so that the host system can access LUNs from those storage systems. (This procedure was performed on an ESS Model 800).
 - e. Be familiar with creating LUNs on the ESS Model 800 so that the host can access the ESS Model 800 devices.
 - f. Although SDD functions correctly in single-path environments, it is recommended that there be redundant physical paths to the devices from the host after installation of SLES 9.
 - g. Optionally, have an understanding of how the Linux kernel boot process functions and what processes and procedures that are used to boot a Linux distribution for a local storage device.
 - h. Ensure that there will be network access to the system.

2. Configure QLogic Devices

- For ease of installation and to avoid issues with internal SCSI or IDE controllers, it is recommended that all internal disk drive controllers be disabled. This procedure assumes that this has been done.
- Verify that the QLogic SAN HBA devices that are configured for the host have been setup to have their BOOT BIOS enabled. This permits discovery and use of SAN disk devices during this procedure. While in the QLogic Utility, configure the ESS Model 800 device from which the system will boot. If the utility cannot see the correct device, check the SAN and ESS Model 800 configurations before continuing.

3. Configure Boot/Root/SWAP devices

The boot device that will be used for installation and booting should be at least 4 GB in size. This is the minimum size for installing a base package set from the installation media to the boot devices.

The swap device be at least the size of physical memory that is configured in the host. For simplicity, these instructions assume that the boot, root, and swap devices are all located on the same device; however, this is not a requirement for the installation.

4. Installation Media

The installation media; that is, the source for installation, can be CD-ROM, NFS, HTTP, FTP, and so forth. For this installation, an NFS-exported set of CD-ROMs was used. You can use any of the installation sources that are listed.

5. Install

- From the Qlogic BIOS Menus, select the installation source to boot from. Verify that the QLogic XXXXXXXX SAN HBA module is loaded and that the SAN devices that will be used for installation have been detected successfully.
- For Emulex fibre HBAs, use Emulex's utility software for the Emulex model to enable the Emulex HBA BIOS (to use this utility, the system must be booted to DOS). After the BIOS is enabled, go into the Emulex BIOS during POST boot and enable the boot BIOS for each adapter and select the boot LUN from the list.
- Because of the way Linux discovers SAN devices, and if SAN devices have already been configured for multiple path access, Linux will discover the same physical device multiple times, once for each logical path to the device. Note which device will be used for the installation before proceeding, for example, /dev/sda.
- Select the desired options until arriving at the Installation Settings. Here, modifications of the partitioning settings are required for this installation. This is to make sure that the device noted in the previous step will be used for the root/boot installation target.
- The details of installation and partitioning are not included here. See the installation procedures to determine which packages are needed for the type of system being installed.

6. Rebooting

- a. On reboot, modify the BIOS to boot from hard disk, the system should now boot to the newly installed OS.
- b. Verify that the system is booted from the correct disk and vpaths.
- c. At this point the installed boot device can be set as the default boot device for the system. This step is not required, but is suggested because it enables unattended reboots after this procedure is complete.

7. Upgrading the SDD driver

At the end of this document are instructions on how to upgrade the SDD driver.

SAN boot configuration

The system should be booted in single-path mode from the SAN. Verify this by looking at the mounted devices and where the root is mounted, using the **df** command. Also verify that the swap, using **swapon -s** and other configured partitions are correctly mounted. This completes the single-path boot from SAN. To modify the boot/root and other devices for booting using the SDD driver continue with the following steps.

Note: The following instructions are examples and the values used herein might be different on your systems. Also, any time a kernel is upgraded a new **initrd** for that kernel must be built following the instructions in the procedure.

1. Install the IBM SDD driver.

Download and install the IBM SDD driver for the kernel version being used. SDD is packaged in an RPM format and can be installed using the **rpm** command. See “Installing SDD” on page 205 for more information.

2. Gather SDD data in preparation for configuring `/etc/fstab`, `menu.lst` and `/boot/initrd`.

```
sdd start
```

The `/etc/vpath.conf` file has now been created. You must ensure that `vpatha` is the root device. Use the **cfgvpath query device** command to obtain the LUN ID of the root's physical device. (In this procedure, `sda` is the root device).

The **cfgvpath query device** command will produce output similar to the following: Note that some data from the following output has been modified for ease of reading.

```
cfgvpath query
/dev/sda (8, 0) host=0 ch=0 id=0 lun=0 vid=IBM pid=2105800 serial=12020870
lun_id=12020870
/dev/sdb (8, 16) host=0 ch=0 id=0 lun=1 vid=IBM pid=2105800 serial=12120870
lun_id=12120870
/dev/sdc (8, 32) host=0 ch=0 id=0 lun=2 vid=IBM pid=2105800 serial=12220870
lun_id=12220870
/dev/sdd (8, 48) host=0 ch=0 id=0 lun=3 vid=IBM pid=2105800 serial=12320870
lun_id=12320870
```

The `lun_id` for `/dev/sda` is 12020870. Edit the `/etc/vpath.conf` file using the `lun_id` for `vpatha`. Remove all other entries from this file (they will be automatically added later by SDD).

3. Modify the `/etc/fstab`, ensuring that `root/boot/swap` is mounted on `vpath` devices.

There is a one-to-one correlation between `sd` and `vpath` minor devices, for example, `sda1` and `vpatha1`.

Major devices, however, might not necessarily correlate, for example, `sdb1` could be `vpathd1`.

Because `/boot` was installed on `/dev/sda1` and `vpatha` corresponds to `sda` in the `/etc/vpath.conf` file, `/dev/vpatha1` will be the mount device for `/boot`.

Example:

Change from:

```

/dev/sda3      /      ext3  defaults  1 1
/dev/sda1      /boot  ext3  defaults  1 2
/dev/sda2      swap   swap  defaults  0 0

```

To:

```

/dev/vpatha3   /      ext3  defaults  1 1
/dev/vpatha1   /boot  ext3  defaults  1 2
/dev/vpatha2   swap   swap  defaults  0 0

```

Modify the `/boot/grub/menu.lst` file. Add entries for the SDD boot using `initrd.vp`

```

title Linux-sdd
kernel (hd0,0)/vmlinuz root=/dev/vpatha3 selinux=0 splash=silent barrier=off
resume=/dev/sda2 elevator=cfq showopts ramdisk_size=34000
initrd (hd0,0)/initrd.vp

```

4. Prepare the `initrd` file.

The `[initrd file]` refers to the current `initrd` in `/boot`. The correct `initrd` can be determined by the following method:

```

ls -lA /boot | grep initrd | grep $(uname -r)
cd /boot
cp [initrd file] to initrd.vp.gz
gunzip initrd.vp.gz
mkdir /boot/mnt

```

For the `ext2` file system, the `initrd` might be required to resize the `initrd` file system.

```

dd if=/dev/zero of=initrd.vp seek=33554432 count=1 bs=1
losetup /dev/loop0 initrd.vp
e2fsck -f /dev/loop0
resize2fs -f /dev/loop0
losetup -d /dev/loop0

```

Adding the `ramdisk_size=` option to the kernel entry in the boot loader file might be required after increasing the size of the `initrd` file. For resizing the `initrd` to 33554432 add the following to the `/boot/grub/menu.lst` file, `ramdisk_size=34000` as mention previously.

5. Change directory to `/boot` and un-archive the `initrd` image to `/boot/mnt`.

Mount the `initrd` file system.

```

mount -o loop -t ext2 initrd.vp /boot/mnt

```

6. Make additional directories in `/boot/mnt` if they do not exist.

```

cd /boot/mnt
mkdir mnt
mkdir -p opt/IBMsdd/bin
chmod -R 640 opt/IBMsdd
mkdir -p lib/tls

```

7. Copy the following `/etc` files to `/boot/mnt/etc`.

```

cp /etc/vpath.conf /boot/mnt/etc/
cp /etc/group /boot/mnt/etc/
cp /etc/passwd /boot/mnt/etc/
cp /etc/nsswitch.conf /boot/mnt/etc/

```

8. Create an fstab file in the initrd etc directory with the following entry (this might already exist).

```
sysfs /sys sysfs defaults 0 0
```

9. Modify the /boot/mnt/etc/nsswitch.conf file.

- a. Change:

```
passwd: compat
```

To:

```
passwd: files
```

- b. Change:

```
group: compat
```

To:

```
group: files
```

10. Copy **cfgvpath** to the initrd image.

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
```

11. Copy required library files for **cfgvpath**.

Use the **ldd** command to determine the library files and locations.

Example:

```
ldd /opt/IBMsdd/bin/cfgvpath | awk {print $(NF-1)} | grep lib
```

The output of the **ldd** command might be similar to the following:

```
/lib/tls/libc.so.6  
/lib/ld-linux.so.2
```

These file must be copied to the /boot/mnt/lib/tls/ and /boot/mnt/lib/ directories respectively.

12. Copy the correct **sdd-mod.o** file to the initrd file system. Use the **uname -r** command to determine the correct **sdd-mod.o** file and create a soft link.

Example:

The **uname -r** command will return something similar to 2.6.5-7.201-smp.

```
cp /opt/IBMsdd/sdd-mod.ko-2.6.5-7.201-smp /boot/mnt/lib/  
cd lib  
ln -s sdd-mod.ko-2.6.5-7.201-smp sdd-mod.ko  
cd ../
```

13. Copy the following binaries to the initrd bin directory.

Note: **mount** and **umount** might already exist. If they do exist, do not copy them to the initrd mount directory.

```
cp /bin/tar /boot/mnt/bin/  
cp /bin/chown /boot/mnt/bin/
```

- Copy the required library files for each binary.

Use the **ldd** command to determine the library files and locations.

Note: Many binaries use the same libraries, so there might be duplications of copying.

Example:

```
ldd /bin/mknod | awk '{print $(NF-1)}' | grep lib
/lib/libselinux.so.1
/lib/tls/libc.so.6
/lib/ld-linux.so.2
```

The above files must be copied to the `/boot/mnt/lib/tls/` and `/boot/mnt/lib/` directories respectively.

Also, copy the following library file to `/boot/mnt/lib/`.

```
cp /lib/libnss_files.so.2 /boot/mnt/lib
```

- Modify the `/boot/mnt/linuxrc` file.

Add the following lines just after the last kernel module is loaded.

```
echo "Loading SDD module"
insmod /lib/sdd-mod.ko
echo "Creating vpath devices"
/opt/IBMsdd/bin/cfgvpath
```

Ensure that an updated copy of `vpath.conf` is copied to the `/root` file system by using the following syntax to mount the root file system.

```
/bin/mount -o rw -t [fstype] [device] /mnt
```

Add the following lines just after the **cfgvpath** command. The values used for the `[fstype]` and `[device]` here are only examples. Use the correct values for the system that is being configured.

```
/bin/mount -o rw -t ext3 /dev/vpatha3 /mnt
/bin/cp /etc/vpath.conf /mnt/etc/
cd /mnt
/bin/tar cps /dev/IBMsdd /dev/vpath* | /bin/tar xps
cd /
/bin/umount /mnt
```

- Prepare the mounted `initrd` file system for rebooting:

```
cd /boot
umount /boot/mnt
gzip initrd.vp
mv initrd.vp.gz initrd.vp
cd /
shutdown -r now
```

- Once booted, verify that `vpath` devices are being used. Add all other paths and reboot again. The following commands can be used to verify the use of `vpath` devices.

```
mount
swapon -s
lsvpcfg
datapath query device
```

SDD upgrade procedure

Use the following procedure to upgrade SDD:

1. Place the SDD rpm update package on the system
2. Restart the system in single-path mode (no SDD)
3. You might come to a **Repair filesystem** prompt. If so, enter the root password. At the prompt, remount the file system as read/write.

```
mount -n -o remount,rw /
```

4. For systems that have /boot on a separate mount point, mount /boot partition using /dev/sd device.
5. Remove the previous SDD driver.

```
rpm -e IBMsdd
```

The /etc/vpath.conf file will be saved to vpath.conf.rpmsave.

6. Install the new SDD driver.

```
rpm -ivh IBMsdd-x.x.x.x-y.i686.sles9.rpm
cd /boot
mv initrd.vp initrd.vp.gz
gunzip initrd.vp.gz
mount -o loop -t ext2 initrd.vp mnt
cp /opt/IBMsdd/sdd-mod.ko-`uname -r` /boot/mnt/lib/
```

7. Verify that the soft link sdd-mod.ko in /boot/mnt/lib points to the current SDD module.
8. Copy the new **cfgvpath** command and use the **ldd** command to verify that the correct libraries are installed for /boot/mnt/opt/IBMsdd/bin/cfgvpath.

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
```

9. Prepare the initrd.vp for booting.

```
umount mnt
gzip initrd.vp
mv initrd.vp.gz initrd.vp
cd /
```

10. Reboot the system.

SAN Boot instructions for SLES 9 with IBM (ppc)

The following procedure is used to install SLES 9 ppc onto a IBM System p JS20 LPAR with fibre-channel connected storage. It is assumed that no installation is present to work from and when completed, the boot and swap devices will be running on IBM SDD vpath devices.

Prerequisite steps

1. Ensure that the following conditions exist before continuing with this procedure:
 - Have a copy of SLES 9 SP2 either network-accessible or on CD-ROM.
 - Be familiar with the SLES installation. This includes understanding which packages will be installed and how to select required options through the installation.

- Be familiar with how to connect to and operate IBM BladeCenter JS20 or IBM System p LPAR.
 - Be familiar with how to set up an LPAR with processors, memory, and SAN HBAs. For network installs a network port is required, and for CD-ROM installs, a CD-ROM is required.
 - Be familiar with how to set up a SAN network or direct-attached SAN storage devices so that the configured system can access LUNs from the storage unit.
 - Be familiar with creating LUNs on the storage unit so that the LPAR can access the storage devices. Although SDD functions correctly in single-path environments, there should be redundant physical paths to the devices from the host (after installation).
 - Optionally, have an understanding of how the Linux kernel boot process functions and what processes and procedures that are used to boot a Linux distribution for a local storage device.
2. Configure root/boot/swap devices
 - The physical boot device that will be used for installation and booting should be at least 4 GB in size. This is the minimum size for installing all packages from the installation media to the boot devices. It is also recommended that the swap device be at least the size of physical memory that is configured in the system. For simplicity, these instructions assume that the root/boot/swap devices are all located on the same device; however this is not a requirement for the installation. Also, it is not required that a /boot mount exists. In some cases, there will not be a /boot mount but rather the boot files will reside in the directory /boot on the root "/" mount.
 3. Installation Media
 - The installation media; that is, the source for installation, can be CD-ROM, NFS, HTTP, FTP, and so forth. For this installation, an NFS-exported set of CD-ROMs was used. Any of the installation sources listed can be used.
 4. Upgrading the SDD driver.

At the end of this document are instructions on how to upgrade the SDD driver.

Use this procedure to install SLES 9:

1. From the SMS menu, select the installation source and boot from the media.
2. Verify that the Emulex lpfcd SAN HBA module is loaded and that the SAN devices that will be used for installation have been detected successfully.

Note: Because of the way Linux discovers SAN devices, and if SAN devices have already been configured for multiple path access, Linux will discover the same physical device multiple times, once for each logical path to the device. Take note which device will be used for the installation before proceeding, that is, /dev/sdh. Also note which of the Emulex devices is used to discover this device as it will be needed in a later step.

3. Select the desired options until arriving at the Installation Settings step of the yast install.

Here, modification of the partitioning settings is required for this installation. This is to make sure that the device noted in the previous step will be used for the root/boot installation target.

 - a. Select partitioning, and go to the custom partition setup.
 - b. Select the device and Custom partitioning for experts.

- c. Make sure that there is a PReP boot partition on the root/boot device and that it is the first partition.
 - d. Continue to partition devices as required to complete this configuration.
The details of installation and partitioning are not written up here. See the installation procedures to determine which packages are needed for the type of system being installed.
4. Finish the installation. An error occurs while attempting to create the yaboot boot loader stating that the device type of *fcv* is unknown. Select **OK** and select **No** when asked to retry.
 5. Rebooting
 - a. On reboot after initial install, enter the SMS menu.
 - b. Boot from the installation source media.
 - c. If you are installing from CD media, continue to a point where you can abort the installation and return to the command line menu system.
 - d. If you are booting from the network, you should already be presented with this menu.
 - e. Select to boot an installed system.
 - f. Select the root device that was just installed in the previous steps. Yast will again come up but from the root partition.
 - g. Finish the installation.
 6. Upgrading to the latest service pack
If there is a service pack available, at the time of this writing there is currently SP2 available, upgrade the installed system to the latest service pack using `yast`. Once this is complete, view `/etc/lilo.conf` and verify that the data in this file looks correct for the boot and root partitions. Once this is verified run `lilo`. This permits the installation of the boot loader to the PReP boot partition of the drive where the installation error occurred from above.
 7. Rebooting.
 - a. Reboot again and enter the SMS menu.
This time the boot device which has been setup over the previous steps is now ready to be booted.
 - b. Select to boot from a Hard Drive/SAN and select the Emulex device adapter associated with the SAN disk device on which the installation was completed.
 - c. The installation boot device should now be listed in the bootable devices discovered in the SAN on the selected Emulex HBA.
 - d. Select the appropriate device and boot.

SAN boot configuration

The system should now be booted in single-path mode from the SAN. Verify this by looking at the mounted devices and where the root is mounted, using the `df` command. Also verify that the swap, using `swapon -s` and other configured partitions are correctly mounted. This completes the single-path boot from SAN. To modify the boot/root and other devices for booting using the SDD driver continue with the following steps.

Note: The following instructions are examples and the values used herein might be different on your systems. In some cases, there will not be a `/boot` mount but rather the boot files will reside in the directory `/boot` on the root `"/"` mounted file system. It is recommended but not required that `vpatha` be used as the `vpath` boot device.

1. Install the IBM SDD driver.

Download and install the IBM SDD driver for the kernel version being used. SDD is packaged in an RPM format and can be installed using the **rpm** command. See “Installing SDD” on page 205 for more information.

2. Extracting the initrd.

The following will unzip and extract the initrd image so that it can be modified to include the required elements to enable a vpath boot image.

```
cd /boot
```

Locate the initrd image used for booting. This will be the image that `/etc/yaboot.conf` is pointing to. Note that the file pointed to might be symbolic link to another file. Copy the file to a temporary filename with a `.gz` extension; that is, if the file name is `initrd-2.6.5-7.191-pseries64` then:

```
cp initrd-2.6.5-7.191-pseries64 initrd.vp.gz
```

Unzip the file using **gunzip**.

```
gunzip initrd.vp.gz
```

Create a temporary directory where the image will be manipulated, for example, `/boot/mnt`. This is referred to as the *image temporary directory* throughout the rest of this documentation. Extract the image to that directory using the command:

```
mkdir -p /boot/mnt
cd /boot/mnt
cpio -iv < ../initrd.vp
```

3. Modifying the `/boot/initrd`.

Create the following directories in the image temporary directory.

For SLES 9 on System p, there might already be a `mnt` directory in the temporary initrd image. If there is not, create one.

```
mkdir mnt
mkdir dev
mkdir -p lib/tls
mkdir -p lib64/tls
mkdir -p opt/IBMsdd/bin
chmod -R 640 opt/IBMsdd
```

Copy the following files to the following directories relative to the image temporary directory.

```
cp /opt/IBMsdd/sdd-mod.ko-2.6.5-7.191-pseries64 lib/
cp /opt/IBMsdd/bin/cfgvpath opt/IBMsdd/bin/
cp /bin/cat bin/
cp /bin/cp bin/
cp /bin/chown bin/
```

For each of the above binary files (except `sdd-mod.o`), run the **ldd** command and verify that the listed library files exist in the image temporary directory. If they do not, copy the listed library files that do not exist to the corresponding `lib` and `lib64` directories in the image temporary directory.

An example script to gather the correct libraries and copy them to the correct directories:

```
for libs in `opt/IBMsdd/bin/cfgvpath /bin/cat /bin/cp /bin/chown`; do
  ldd $libs | awk '{print $(NF-1)}' | grep lib |
  while read line; do
    cp $line /boot/mnt$line
  done
done
```

Copy this additional library file.

4. Gather SDD Data in preparation for configuring `/etc/fstab`, `/etc/yaboot.conf` and `/boot/initrd`.

```
sdd start
```

The `/etc/vpath.conf` file has now been created. You must ensure that `vpatha` is the root device. Use the **cfgvpath query device** command to obtain the LUN ID of the root's physical device. (In this procedure, `sda` is the root device). The **cfgvpath query** command produces output similar to the following example.

Note that some data from the following output has been modified for ease of reading.

```
cfgvpath query

/dev/sda (8, 0) host=0 ch=0 id=0 lun=0 vid=IBM pid=2105800 serial=12020870
  lun_id=12020870
/dev/sdb (8, 16) host=0 ch=0 id=0 lun=1 vid=IBM pid=2105800 serial=12120870
  lun_id=12120870
/dev/sdc (8, 32) host=0 ch=0 id=0 lun=2 vid=IBM pid=2105800 serial=12220870
  lun_id=12220870
/dev/sdd (8, 48) host=0 ch=0 id=0 lun=3 vid=IBM pid=2105800 serial=12320870
  lun_id=12320870
```

The `lun_id` for `/dev/sda` is 12020870. Edit the `/etc/vpath.conf` file using the `lun_id` for `vpatha` (`vpatha 12020870`). Remove all other entries from this file (they will be automatically added later by SDD).

5. Modify the `/etc/fstab` to use `vpath` devices for `/root` and `swap`.

Other devices using `vpaths` will also need to be changed. For the initial install, work only with the `root/boot/swap` devices and comment out other `sd` and `hd` devices until completed.

Original:

```
/dev/sdd4      /          ext3      acl,user_xattr  1 1
/dev/hda2     /data1    auto      noauto,user     0 0
/dev/hda4     /data2    auto      noauto,user     0 0
/dev/hda3     swap      swap      pri=42          0 0
/dev/sdd3     swap      swap      pri=42          0 0
devpts       /dev/pts  devpts    mode=0620,gid=5 0 0
proc         /proc     proc      defaults        0 0
usbfs        /proc/bus/usb usbfs     noauto          0 0
sysfs        /sys      sysfs     noauto          0 0
```

Modified:

```
/dev/vpatha4  /          ext3      acl,user_xattr  1 1
#/dev/hda2   /data1    auto      noauto,user     0 0
#/dev/hda4   /data2    auto      noauto,user     0 0
#/dev/hda3   swap      swap      pri=42          0 0
/dev/vpatha3 swap      swap      pri=42          0 0
devpts       /dev/pts  devpts    mode=0620,gid=5 0 0
proc         /proc     proc      defaults        0 0
usbfs        /proc/bus/usb usbfs     noauto          0 0
sysfs        /sys      sysfs     noauto          0 0
```

6. Stop SDD and copy /etc files to the image temporary directories.

```
sdd stop
cp /etc/vpath.conf /boot/mnt/etc
cp /etc/passwd /boot/mnt/etc
cp /etc/group /boot/mnt/etc
```

7. Edit the /boot/mnt/linuxrc file.

Edit the init file in the image temporary directory. Go to the line that has the creating device nodes message, right after the init script creates the device nodes using /bin/udevstart, and add the following lines after the echo -n . command in the script block.

```
echo "Creating vpath devices"
/opt/IBMsdd/bin/cfgvpath

echo "Mounting and copying some required SDD files"
/bin/mount -o rw -t <PARTITION TYPE> /dev/vpathXXX /mnt

/bin/cp /etc/vpath.conf /mnt/etc
/bin/umount /mnt
insmod /lib/scsi_mod.o
```

where /dev/vpathXXX is the root drive/partition.

8. Edit the /boot/mnt/load_modules.sh file.

Edit the load_modules.sh file in the image temporary directory and add the following lines to the end of the script:

```
echo "Loading SDD Module"
insmod /lib/sdd-mod.ko
```

Issue the following command from the lib directory. The linked module is the name of the module that was copied into the lib directory above.

```
cd /boot/mnt/lib
ln -s sdd-mod.ko-2.6.5-7.191-pseries64 sdd-mod.ko
```

9. Repackaging the initrd.

To repack all of the changes that have just been made to the initrd, issue the following commands:

```
cd /boot/mnt
find . | cpio -H newc -vo > /boot/initrd.vp
cd /boot
gzip initrd.vp
mv initrd.vp.gz initrd.vp
```

The initrd-2.6.5-7.191-pseries64 now has the repackaged initrd image with the SDD drive and modified files required to boot from a vpath.

10. Modifying root device files.

Additional files need to be modified on the root file system before the modifications are complete.

Modify /etc/yaboot.conf. Add a new entry in the file and modify the entry to point at the new initrd image created in the above step. Also modify the root device in the new entry to point to the vpath chosen from the previous steps. Remember to include the partition if required. Also make sure to modify the entry name.

Original /etc/yaboot.conf:

```
# header section
partition = 4
timeout = 100
default = linux
# image section
image = /boot/vmlinux
label = linux
append = "root=/dev/sdd4 selinux=0 elevator=cfq"
initrd = /boot/initrd
```

Modified /etc/yaboot.conf:

```
# header section
partition = 4
timeout = 100
default = linux
# image section
image = /boot/vmlinux
label = linux
append = "root=/dev/sdd4 selinux=0 elevator=cfq"
initrd = /initrd
image = /boot/vmlinux
label = linux-sdd
append = "root=/dev/vpatha3 selinux=0 elevator=cfq"
initrd = /boot/initrd.vp
```

11. Restart the system.
 - a. Reboot the system.
 - b. From the SMS menu, select the boot devices as before, if the boot device is not already set up as the first boot device.
 - c. When the yaboot prompt is shown during boot, enter the name for the new boot image.
 - d. During the OS load, ensure that the IBMsdd module is loaded after the SAN disk devices are discovered.
 - e. Ensure that no errors were printed to the console during boot.
 - f. If there were errors, reboot the system and at the yaboot prompt, select the old image to boot from. When the system boots, review the previous steps and correct any errors. Then repeat these steps, starting with step 9 (repackaging the initrd).

If all the vpath devices that are discovered by **cfgvpath** were not created during the modification steps above, the **cfgvpath** might have to timeout waiting for these devices to be created.

Once the system comes up, login and verify that the root mount device is the device specified during the configuration, using **df**. Also validate that any other configured partitions and swap devices, using **swapon -s**, are also now mounted on vpath devices.

12. Verify System has reboot and SDD is configured correctly.

Once booted, verify that vpath devices are being used. Add all other paths and reboot again. The following commands can be used to verify the use of vpath devices:

- **mount**
- **swapon -s**
- **lsvpcfg**
- **datapath query device**

At this point, the installed boot device can be set as the default boot device for the system. This step is not required, but is suggested because it enables unattended reboots after this procedure is complete.

Upgrading the SDD

Use the following procedure to upgrade SDD:

1. Place the SDD rpm update package on the system.
2. Restart the system in single-path mode (no SDD).
3. You might come to a **Repair filesystem** prompt. If so, enter the root password. At the prompt, remount the file system as read/write.

```
mount -n -o remount,rw /
```

For systems that have /boot on a separate mount point, mount /boot partition using /dev/sd device.

4. Remove the previous SDD driver.

```
rpm -e IBMsdd
```

The /etc/vpath.conf file will be saved to vpath.conf.rpmsave.

5. Install the new SDD driver.

```
rpm -ivh IBMsdd-x.x.x.x-y.ppc64.sles9.rpm
mkdir -p /boot/mnt
cd /boot
mv initrd.vp initrd.vp.gz
gunzip initrd.vp.gz
cd /boot/mnt
cpio -iv < ../initrd.vp
cp /opt/IBMsdd/sdd-mod.ko-`uname -r` /boot/mnt/lib
```

6. Verify that the soft link sdd-mod.ko in /boot/mnt/lib points to the current SDD module.
7. Copy the new **cfgvpath** command and use the **ldd** command to verify that the correct libraries are installed for /boot/mnt/opt/IBMsdd/bin/cfgvpath.

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
```

8. Prepare the initrd.vp for booting.

```
cd /boot/mnt
find . | cpio -H newc -vo > /boot/initrd.vp
cd /boot
gzip initrd.vp
mv initrd.vp.gz initrd.vp
rm -rf mnt
cd /
```

9. Reboot the system.

SAN Boot Instructions for SLES 9 with IBM SDD (x86) and LVM 2

The following procedure is used to install SLES 9 x86 on an xSeries host with fibre-channel connect storage and configure SDD with LVM. This procedure assumes that no installation is present to work from and when completed, the boot and swap devices will be running on IBM SDD vpath devices and will be under LVM control.

Prerequisite steps

1. Ensure that the following conditions exist before continuing with this procedure:
 - a. The installation target **MUST** be single-pathed before installing SLES 9. It is also recommended to limit the installation to a single LUN if possible to ease the transition from single-path to IBM SDD vpath; however, this is not required.
 - b. The QLogic BIOS should be enabled for discovery of SAN Devices and that the device that contains the kernel and initrd images, /boot mount point, be selected as the boot device in the QLogic BIOS. Follow the *IBM Host Systems Attachment Guide* recommendations when setting up the QLogic BIOS for SAN Boot.
 - c. Have a copy of SLES 9 SP2 i386 either network-accessible or on CD-ROM.
 - d. Be familiar with the SLES 9 installation. This includes understanding which packages will be installed.
 - e. Be familiar with setting up root, boot, swap and any other initial mount points that will be used for the setup of the initial system under LVM control.
 - f. Be familiar with how to set up a SAN network or direct-attached SAN storage devices so that the host system can access LUNs from those storage systems.
 - g. Be familiar with how to set up a SAN network or direct-attached SAN storage devices so that the host system can access LUNs from the those storage systems. Although SDD functions correctly in single-path environments, it is recommended that there be redundant physical paths to the devices from the host after completing this procedure.
 - h. Optionally, have an understanding of how the Linux kernel boot process functions and what processes and procedures that are used to boot a Linux distribution for a local storage device.
 - i. Ensure that there will be network access to the system.
2. Configure QLogic Devices.
 - For ease of installation and to avoid issues with internal SCSI or IDE controllers, it is recommended that all internal disk drive controllers be disabled. This procedure assumes that this has been done.
 - Verify that the QLogic SAN HBA devices that are configured for the host have been set up to have their BOOT BIOS enabled. This permits discovery and use of SAN disk devices during this procedure. While in the QLogic Utility, configure the ESS Model 800 device from which the system will boot. If the utility cannot see the correct device, check the SAN and ESS Model 800 configurations before continuing.
3. Configure Boot/Root/SWAP devices.

The root device that will be used for installation and booting should be at least 4 GB in size. If multiple partitions are being used, that is, /usr /var, the total size of all mount points should be at least this size. This is the minimum size for installing a base package set from the installation media to the boot devices. More space might be required depending on the package selection.

The swap device be at least the size of physical memory that is configured in the host. For simplicity, these instructions assume that the boot, root, and swap devices are all located on the same device; however, this is not a requirement for the installation.

 - The boot (/boot) device must **NOT** be under LVM control.

- The root, (/), and other optional mount points, (/usr, /var, /opt), can be under LVM control. If they are not, at a minimum, they should be mounted to an IBM SDD vpath device.
 - SWAP can also be under LVM control but this is not a requirement but should at least use a vpath device.
4. Use the installation media.
The installation media; that is, the source for installation, can be CD-ROM, NFS, HTTP, FTP, and so forth. For this installation, an NFS-exported set of CD-ROMs was used. You can use any of the installation sources that are listed.
 5. Installing the system.
 - From the BIOS Menu, select the installation source to boot from. Verify that the QLogic qla2300 SAN HBA module is loaded and that the SAN devices that will be used for installation have been detected successfully.
 - Because of the way Linux discovers SAN devices, and if SAN devices have already been configured for multiple path access, Linux will discover the same physical device multiple times, once for each logical path to the device. Note which device will be used for the installation before proceeding, for example, /dev/sda.
 - Select the desired options until arriving at the Installation Settings. Here, modifications of the partitioning settings are required for this installation. This is to make sure that the device noted in the previous step will be used for the root/boot installation target.
 - The details of installation, partitioning, LVM setup, package selection, boot options, and so on, are not documented here. See the installation procedures to determine which packages are needed for the type of system that you are installing.
 6. Restarting the system.
 - a. On reboot, modify the BIOS to boot from hard disk, the system should now boot to the newly installed OS.
 - b. Verify that the system is booted from the correct disk and vpaths.
 - c. At this point the installed boot device can be set as the default boot device for the system. This step is not required, but is suggested because it enables unattended reboots after this procedure is complete.

SAN boot configuration

The system should be booted in single-path mode from the SAN. Verify this by looking at the mounted devices and where the root is mounted, using the **df** command. Also verify that the swap, using **swapon -s** and other configured partitions are correctly mounted. This completes the single-path boot from SAN. To modify the boot/root and other devices for booting using the SDD driver continue with the following steps.

The following list contains notes about the install procedure that should be reviewed before beginning:

- All values and devices in the following procedure might not be the same on the system where this procedure is being conducted. It is, however, recommended (but not required) that you use *vpatha* as the physical device for the root volume group.
- Perform this procedure in a single-path environment. Once completed and booting with SDD and LVM, configure the SAN for multipath.
- All commands in this procedure begin with a # sign and might be followed by the output of that command such as the command **pvdiskdisplay**.

- Because /boot will not be under LVM control, it might be safer to work from within /boot. - I
- In this procedure, you will work with a copy of the current initrd named initrd.vp.
- The volume groups for root and swap in the example are as follows:

```

/dev/rootVolGroup/
/dev/rootVolGroup/rootVol
/dev/rootVolGroup/swapVol
/dev/rootVolGroup/rootVol -> /dev/mapper/rootVolGroup-rootVol
/dev/rootVolGroup/swapVol -> /dev/mapper/rootVolGroup-swapVol
Physical device is sda2
vpath device vpatha2

```

- Before starting SDD, comment out any sd devices from /etc/fstab other than /boot. This will ensure all devices are written to the /etc/vpath.conf file. These devices might later be changed to vpath devices if the intent is to have them multipathed.
- The /etc/fstab will also need to be modified to point /boot from /dev/sd[x] or LABEL=[some_label_name_here] to /dev/vpath[x].
- Modify the /boot/grub/menu.lst file to add an entry for the SDD initrd.
- Modify /etc/lvm/lvm.conf to recognize vpath devices and ignore sd devices.
- It is always a good idea to make copies of files that are going to be manually modified such as /etc/fstab, /etc/vpath.conf /etc/lvm/lvm.conf and /boot/grub/menu.lst.

1. Install the IBM SDD driver.

Download and install the IBM SDD driver for the kernel version being used. SDD is packaged in an RPM format and can be installed using the **rpm** command. See “Installing SDD” on page 205 for more information.

2. Use **pvdisplay** to show the physical volume(2) currently configured for use in LVM. These volumes will be converted from a single-path sd drive to the IBM SDD vpath device. The following is an example of the output from **pvdisplay**.

```

# pvdisplay

--- Physical volume ---
PV Name           /dev/sda2
VG Name           rootVolGroup
PV Size           9.09 GB / not usable 0
Allocatable       yes
PE Size (KByte)   32768
Total PE          291
Free PE           1
Allocated PE      290
PV UUID           SSm5g6-UoWj-evHE-kBj1-3QB4-EVi9-v88xiI

```

3. Modify the /etc/fstab, ensuring that:

- a. LABEL= is not being used.
- b. /boot is mounted on a vpath device.

There is a one-to-one correlation between sd and vpath minor devices, such as, sda1 and vpatha1.

Major devices, however, might not necessarily correlate; for example, sdb1 could be vpathd1.

Because /boot was installed on /dev/sda1 and vpatha corresponds to sda in the/etc/vpath.conf file,/dev/vpatha1 will be the mount device for /boot.

Example:

Change from:

```
/dev/rootVolGroup/rootVol /      ext3 defaults 1 1
LABEL=/boot                /boot  ext3 defaults 1 2
/dev/rootVolGroup/swapVol  swap   swap  defaults 0 0
```

To:

```
/dev/rootVolGroup/rootVol /      ext3 defaults 1 1
/dev/vpath1                /boot  ext3 defaults 1 2
/dev/rootVolGroup/swapVol  swap   swap  defaults 0 0
```

4. Modify the `/boot/grub/menu.lst` file. Add an entry before the first title entry for the SDD/LVM boot using `initrd.vp`. Verify which is the default boot image. The default line should point to the new entry. Make sure the root and resume are identical to the current Linux installation.

```
...
title Linux w/LVM w/SDD
kernel (hd0,0)/vmlinuz root=/dev/system/lv00 resume=/dev/system/swap
selinux=0 splash=silent barrier=off elevator=cfq
initrd (hd0,0)/initrd.vp
...
```

5. Modify `/etc/lvm/lvm.conf`.

This procedure will modify LVM to only discover vpath style devices.

Comment out the default filter line.

Example:

```
filter = [ "a./" ]
```

Add the following line:

```
filter = [ "a/vpath*", "r/sd*" ]
```

In the `types =` section add an entry for vpath

```
types = [ "vpath", 16 ]
```

6. Modify the boot scripts.

To support the addition of vpath devices during boot as well as possible changes to the device-mapper, you must add and modify the following boot scripts.

```
# cd /etc/init.d/boot.d
# ln -s ../boot.udev S04boot.udev
# vi S06boot.device-mapper
```

Add the following lines after `/sbin/devmap_mknod.sh`:

```
/sbin/dmsetup mknodes

# cd /etc/init.d
# vi boot.local
```

Add the following line to the bottom of the file:

```
/opt/IBMsdd/bin/cfgvpath
```

7. Start the SDD:

```
sdd start
```

The `/etc/vpath.conf` file has now been created. You must ensure that `vpatha` is the root device. Use the `cfgvpath query device` command to obtain the LUN ID of the root's physical device. (In this procedure, `sda` is the root device).

The `cfgvpath query` command produces output similar to the following example. Note that some data from the following output has been modified for ease of reading.

```
# cfgvpath query
/dev/sda (8, 0) host=0 ch=0 id=0 lun=0 vid=IBM pid=2105800 serial=12020870
lun_id=12020870
```

The `lun_id` for `/dev/sda` is 12020870. This is the `sd` device that you will map to `vpatha`. Edit the `/etc/vpath.conf` file using the `lun_id` for `vpatha` and remove all other entries from this file. (SDD will automatically add them later.)

```
vpatha 12020870
```

8. Prepare the `initrd` file.

The [initrd file] refers to the current `initrd` in `/boot`. The correct `initrd` can be determined by the following method:

```
# ls -lA /boot | grep initrd | grep $(uname -r)
cd /boot
cp [initrd file] to initrd.vp.gz
gunzip initrd.vp.gz
mkdir /boot/mnt
```

`initrd-2.6.5-7.191-smp` might be the result.

```
cd /boot
cp [initrd file] to initrd.vp.gz
gunzip initrd.vp.gz
mkdir /boot/mnt
```

For the `ext2` file system, the `initrd` might be required to resize the `initrd` file system.

9. Resize and mount the `initrd` image.

For `x86`-based systems, the `initrd` is an `ext2` filesystem. Because of the need to add files to the `initrd` image, you should increase the size of the image before continuing. After issuing the `e2fsck -f initrd.vp` command you are prompted create a `/lost+found` directory. Enter `y` to create this directory.

```
e2fsck 1.36 (05-Feb-2005)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
/lost+found not found. Create<y>?

# dd if=/dev/zero of=initrd.vp seek=33554432 count=1 bs=1
# e2fsck -f initrd.vp
# resize2fs -f initrd.vp
# mount -o loop -t ext2 initrd.vp /boot/mnt
```

Note: For the remainder of this procedure, work from `/boot/mnt`.

10. Make additional directories in `/boot/mnt` if they do not exist.

```
# mkdir /boot/mnt/mnt
# mkdir -p /boot/mnt/opt/IBMsdd/bin
# chmod -R 640 /boot/mnt/opt/IBMsdd
# mkdir -p /boot/mnt/lib/tls
```

11. Copy the following `/etc` files to `/boot/mnt/etc`.

```
# cp /etc/vpath.conf /boot/mnt/etc/
# cp /etc/group /boot/mnt/etc/
# cp /etc/passwd /boot/mnt/etc/
# cp /etc/nsswitch.conf /boot/mnt/etc/
# cp /etc/fstab /boot/mnt/etc/
# cp /etc/lvm/lvm.conf /boot/mnt/etc/lvm/
```

12. Modify the `/boot/mnt/etc/fstab` file.

Remove all lines that the begin with `/dev/*`

13. Modify the `/boot/mnt/etc/nsswitch.conf` file.

- a. Change:

```
passwd: compat
```

To:

```
passwd: files
```

- b. Change:

```
group: compat
```

To:

```
group: files
```

14. Copy `/opt/IBMsdd/bin/cfgvpath` to `/boot/mnt/opt/IBMsdd/bin/`.

```
# cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
# chmod 755 /boot/mnt/opt/IBMsdd/bin/*
```

15. Copy the required library files for **cfgvpath**.

Use the **ldd** command to determine the library files and locations.

Example:

```
ldd /opt/IBMsdd/bin/cfgvpath | awk {print $(NF-1)}
```

The output of the **ldd** command might be similar to the following:

```
/lib/tls/libc.so.6
/lib/ld-linux.so.2
```

These files must be copied to the `/boot/mnt/lib/tls/` and `/boot/mnt/lib/` directories respectively.

16. Copy the correct `sdd-mod.o` file to the `initrd` file system. Use the **uname -r** command to determine the correct `sdd-mod.o` file and create a soft link.

Example:

The **uname -r** command will return something similar to `2.6.5-7.191-smp`.

```
# cp /opt/IBMsdd/sdd-mod.ko-<uname -r> /boot/mnt/lib/sdd-mod.ko
```

17. Verify that the following files exist in the initrd image. If they do not exist, copy the following binaries:

```
# cp /bin/tar /boot/mnt/bin/  
# cp /bin/awk /boot/mnt/bin/  
# cp /bin/chown /boot/mnt/bin/  
# cp /bin/grep /boot/mnt/bin/  
# cp /bin/mknod /boot/mnt/bin/  
# cp /bin/cp /boot/mnt/bin/
```

18. Copy the required library files for each binary that you copied over in step 15. Use the **ldd** command to determine the library files and locations.

Note: Many binaries use the same libraries, so there might be duplications of copying. Verify that the libraries do not already exist in `/boot/mnt/lib`, if they already exist, there is no need to copy over a new version.

Example:

```
# ldd /bin/mknod | awk '{print $(NF-1)}' | grep lib  
  
/lib/libselinux.so.1  
/lib/tls/libc.so.6  
/lib/ld-linux.so.2
```

The above files must be copied to the `/boot/mnt/lib/tls/` and `/boot/mnt/lib/` directories respectively.

Also, copy the following library file to `/boot/mnt/lib/`.

```
cp /lib/libnss_files.so.2 /boot/mnt/lib
```

19. Modify the `/boot/mnt/linuxrc` file.

Add the following lines just before the statement, `echo Loading kernel/drivers/md/dm-snapshot.ko`.

```
echo "Loading SDD module"  
insmod /lib/sdd-mod.ko  
echo "Creating vpath devices"  
/opt/IBMsdd/bin/cfgvpath
```

Ensure that an updated copy of `vpath.conf` is copied to the `/root` file system by using the following syntax to mount the root file system.

```
/bin/mount -o rw -t [fstype] [device] /mnt
```

Add the following lines just after `[vgchange <...>]`. The values used for the `[fstype]` and `[device]` here are only examples. Use the correct values for the system that is being configured.

```
/bin/mount -o rw -t ext3 /dev/vpatha3 /mnt  
/bin/cp /etc/vpath.conf /mnt/etc/  
cd /mnt
```

20. Prepare the mounted initrd file system for rebooting:

```
# cd /boot
# umount /boot/mnt
# gzip initrd.vp
# mv initrd.vp.gz initrd.vp
# cd /
```

21. Once booted, verify that vpath devices are being used. Add all other paths and reboot again. The following commands can be used to verify the use of vpath devices.

```
mount
swapon -s
pvdisplay
lsvpcfg
datapath query device
```

SAN boot instructions for RHEL 4 with IBM SDD (x86)

The following procedure is used to install RHEL 4 x86 on an xSeries host with fibre-channel connect storage. This procedure assumes that no installation is present to work from and when completed, the boot and swap devices will be running on IBM SDD vpath devices.

Prerequisite steps

1. Ensure that the following conditions exist before continuing with this procedure:
 - a. The installation target MUST be single-pathed before installing RHEL 4.
 - b. Have a copy of RHEL 4 U1 or U2 i386 either network-accessible or on CD-ROM.
 - c. Be familiar with the RHEL 4 installation. This includes understanding which packages will be installed.
 - d. Be familiar with how to set up a SAN network or direct-attached SAN storage devices so that the host system can access LUNs from those storage systems. (This procedure was performed on an ESS Model 800).
 - e. Be familiar with creating LUNs on the ESS Model 800 so that the host can access the ESS Model 800 devices.
 - f. Although SDD functions correctly in single-path environments, it is recommended that there be redundant physical paths to the devices from the host after installation of RHEL 4.
 - g. Optionally, have an understanding of how the Linux kernel boot process functions and what processes and procedures that are used to boot a Linux distribution for a local storage device.
 - h. Ensure that there will be network access to the system.
2. Configure QLogic Devices

Note: For ease of installation and to avoid issues with internal SCSI or IDE controllers, all internal disk drive controllers should be disabled. This procedure assumes that this has been done.

- Verify that the QLogic SAN HBA devices that are configured for the host have been setup to have their BOOT BIOS enabled. This permits discovery and use of SAN disk devices during this procedure. While in the QLogic Utility, configure the ESS Model 800 device from which the system will boot. If the utility cannot see the correct device, check the SAN and ESS Model 800 configurations before continuing.

3. Configure Boot/Root/SWAP devices.

- The boot device that will be used for installation and booting should be at least 4 GB in size. This is the minimum size for installing a base package set from the installation media to the boot devices.
- It is also recommended that the swap device be at least the size of physical memory that is configured in the host. For simplicity these instructions assume that the boot, root, and swap devices are all located on the same device. However, this is not a requirement for the installation.

4. Installation Media

The installation media; that is, the source for installation, can be CD-ROM, NFS, HTTP, FTP, and so forth. For this installation, an NFS-exported set of CD-ROMs was used. You can use any of the installation sources that are listed.

5. Install

- Verify that the QLogic qla2030 SAN HBA module is loaded and that the SAN devices that will be used for installation have been detected successfully.
- For Emulex fibre HBAs, use Emulex utility software for the Emulex model to enable the Emulex HBA BIOS (to use this utility, the system must be booted to DOS). After the BIOS is enabled go into the Emulex BIOS during POST boot and enable the boot BIOS for each adapter and select the boot LUN from the list.
- Because of the way Linux discovers SAN devices, and if SAN devices have already been configured for multiple path access, Linux will discover the same physical device multiple times, once for each logical path to the device. Note which device will be used for the installation before proceeding, that is, /dev/sda.
- Select the desired options until arriving at the Installation Settings. Here, modifications of the partitioning settings are required for this installation. This is to make sure that the device noted in the previous step will be used for the root/boot installation target.
- The details of installation and partitioning are not written up here. See the installation procedures to determine which packages are needed for the type of system being installed.

6. Rebooting

- a. On reboot, modify the BIOS to boot from hard disk; the system should now boot to the newly installed OS.
- b. At this point the installed boot device can be set as the default boot device for the system. This step is not required, but is suggested because it enables unattended reboots after this procedure is complete.

SAN boot configuration

The system should be booted in single-path mode from the SAN. Verify this by looking at the mounted devices and where the root is mounted, using the **df** command. Also verify that the swap, using **swapon -s** and other configured partitions are correctly mounted. This completes the single-path boot from SAN.

The following list of suggestions should be noted before beginning this procedure:

Note:

1. The following instructions are examples and the values used herein might be different on your system. In some cases, there will not be a /boot mount but rather the boot files will reside in the directory /boot on the root "/" mounted file system. It is recommended, but not required, that vpatha be used as the vpath boot device.

2. All values and devices in the following procedure might not be the same on the system where this procedure is being conducted. It is, however, recommended (but not required) to use *vpatha* as the physical device for the root volume group.
3. Perform this procedure in a single-path environment. Once completed and booting with SDD and LVM, configure the SAN for multipath.
4. All commands in this procedure begin with a # sign and might be followed by the output of that command.
5. In this procedure, you will work with a copy of the current `initrd` named `initrd.vp`.
6. Before you start SDD, comment out any `sd` devices from `/etc/fstab` other than `/boot`. This ensures that all devices are written to the `/etc/vpath.conf` file. These devices might later be changed to `vpath` devices if the intent is to have them multipathed.
7. The `/etc/fstab` will also need to be modified to point `/boot` from `/dev/sd[x]` or `LABEL=[some_label_name_here]` to `/dev/vpath[x]`.
8. Modify the `/boot/grub/menu.lst` file to add an entry for the SDD `initrd`.
9. It is always a good idea to make copies of files that are going to be manually modified such as `/etc/fstab`, `/etc/vpath.conf`, `/etc/lvm/lvm.conf` and `/boot/grub/menu.lst`.

To modify the boot/root and other devices for booting using the SDD driver, continue with the following steps:

1. Install the IBM SDD driver.

Download and install the IBM SDD driver for the kernel version being used. SDD is packaged in an RPM format and can be installed using the `rpm` command. See "Installing SDD" on page 205 for more information.

2. Modify the `/etc/fstab` file, ensuring that:
 - a. `LABEL=` is not being used
 - b. `/boot` is mounted on a `vpath` device

Because Red Hat writes labels to the disk and uses labels in the `/etc/fstab` the `/boot` (`/boot`) device might be specified as a label, that is, `LABEL=/boot`. This might, however, be a different label other than `LABEL=/boot`. Check for line in the `/etc/fstab` where `/boot` is mounted and change it to the correct `vpath` device. Also ensure that any other device specified with the `LABEL=` feature is changed to a `/dev/sd` or `/dev/vpath` device. `LABEL=` in a multi-pathed environment confuses Red Hat.

There is a one-to-one correlation between `sd` and `vpath` minor devices, such as `sda1` and `vpatha1`. Major devices, however, might not necessarily correlate; for example, `sdb1` could be `vpathd1`. Because `/boot` was installed on `/dev/sda1` and `vpatha` corresponds to `sda` in the `/etc/vpath.conf` file, `/dev/vpatha1` will be the mount device for `/boot`.

3. Gather SDD data in preparation for configuring `/etc/fstab`, `menu.lst` and `/boot/initrd`.

```
sdd start
```

The `/etc/vpath.conf` file has now been created. You must ensure that `vpatha` is the root device. Use the `cfgvpath query device` command to obtain the LUN ID of the root's physical device. (In this procedure, `sda` is the root device).

The `cfgvpath query` command produces output similar to the following example. Note that some data from the following output has been modified for ease of reading.

```
cfgvpath query
/dev/sda (8, 0) host=0 ch=0 id=0 lun=0 vid=IBM pid=2105800 serial=12020870
lun_id=12020870
/dev/sdb (8, 16) host=0 ch=0 id=0 lun=1 vid=IBM pid=2105800 serial=12120870
lun_id=12120870
/dev/sdc (8, 32) host=0 ch=0 id=0 lun=2 vid=IBM pid=2105800 serial=12220870
lun_id=12220870
/dev/sdd (8, 48) host=0 ch=0 id=0 lun=3 vid=IBM pid=2105800 serial=12320870
lun_id=12320870
```

The `lun_id` for `/dev/sda` is 12020870. Edit the `/etc/vpath.conf` file using the `lun_id` for `vpatha`. Remove all other entries from this file (they will be automatically added later by SDD).

Add an entry for the SDD/LVM boot using `initrd.vp`

4. Modify the `/boot/grub/menu.lst` file

```
title Red Hat Enterprise Linux AS (2.6.9-11.ELsmp) w/SDD
root (hd0,0)
kernel /vmlinuz-2.6.9-11.ELsmp ro root=/dev/vpatha3
initrd /initrd.vp
```

5. Prepare the `initrd` file.

The `[initrd file]` refers to the current `initrd` in `/boot`. The correct `initrd` can be determined by the following method:

```
ls -lA /boot | grep initrd | grep $(uname -r)

initrd-2.6.9-11.ELsmp.img might be the result.

cd /boot
cp [initrd file] to initrd.vp.gz
gunzip initrd.vp.gz
mkdir /boot/mnt
```

6. Change directory to `/boot` and un-archive the `initrd` image to `/boot/mnt`.

```
cd /boot/mnt
cpio -iv < ../initrd.vp
```

7. Make additional directories in `/boot/mnt`.

```
mkdir mnt
mkdir -p opt/IBMsdd/bin
chmod -R 640 opt/IBMsdd
mkdir -p lib/tls
```

8. Copy the following `/etc` files to `/boot/mnt/etc`.

```
cp /etc/vpath.conf /boot/mnt/etc/
cp /etc/group /boot/mnt/etc/
cp /etc/passwd /boot/mnt/etc/
cp /etc/nsswitch.conf /boot/mnt/etc/
cp /etc/fstab /boot/mnt/etc/
```

9. Copy `/opt/IBMsdd/bin/cfgvpath` to `/boot/mnt/opt/IBMsdd/bin/`

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin
```

- Copy required library files for **cfgvpath**. Use the **ldd** command to determine the library files and locations.

Example:

```
ldd /opt/IBMsdd/bin/cfgvpath | awk '{print $(NF-1)}'
```

The **ldd** command returns something similar to the following:

```
/lib/tls/libc.so.6  
/lib/ld-linux.so.2
```

These files must be copied to the **/boot/mnt/lib/tls/** and **/boot/mnt/lib/** directories respectively.

- Copy the correct **sdd-mod** to the **initrd** file system.

Use the **uname -r** command to determine the correct **sdd-mod** and create a soft link.

Example:

The command will return something similar to **2.6.9-11.ELsmp**

```
cp /opt/IBMsdd/sdd-mod.ko-2.6.9-11.ELsmp /boot/mnt/lib/sdd-mod.ko
```

- Copy the following binaries and change permissions to **755**:

```
cp /bin/mknod /boot/mnt/bin/  
cp /bin/cp /boot/mnt/bin/
```

- Copy required library files for each binary copied to the **/boot/mnt** directory in the previous step. Use the **ldd** command to determine the library files and locations.

Note: Many binaries use the same libraries so there might be duplications of copying. Also, copy the following libraries:

```
ldd /bin/mknod | awk '{print $(NF-1)}' | grep lib  
  
/lib/libselinux.so.1  
/lib/tls/libc.so.6  
/lib/ld-linux.so.2
```

The above files must be copied to the **/boot/mnt/lib/tls/** and **/boot/mnt/lib/** directories respectively. Also, copy the following library files to **/boot/mnt/lib/**.

```
cp /lib/libproc-3.2.3.so /boot/mnt/lib/  
cp /lib/libtermcap.so.2 /boot/mnt/lib/  
cp /lib/libnss_files.so.2 /boot/mnt/lib/
```

- Modify the **/boot/mnt/init** file.

Add the following lines after the modules load and just before the **/sbin/udevstart**. Note that **/sbin/udevstart** can exist multiple times in the **initrd**. Make sure these lines are added before the correct **/sbin/udevstart** entry which is located after the kernel modules load.

```
echo "Loading SDD module"  
insmod /lib/sdd-mod.ko  
echo "Creating vpath devices"  
/opt/IBMsdd/bin/cfgvpath
```

Ensure that an updated copy of `vpath.conf` is copied to the `/root` file system at boot time by using the following syntax to mount the root file system.

```
/bin/mount -o rw -t [fstype] [device] /mnt
```

Add the following lines to the init file just after the previously added entries. Values used for the `[fstype]` and `[device]` here are only examples. Use the correct values for the system that is being configured.

```
/bin/mount -o rw -t ext3 /dev/vpatha3 /mnt  
/bin/cp /etc/vpath.conf /mnt/etc/  
/bin/umount /mnt
```

15. Use **cpio** to archive the `/boot/mnt` directory and **gzip** it in preparation for rebooting.

```
find . | cpio -H newc -vo > ../initrd.vp  
cd /boot  
gzip initrd.vp  
mv initrd.vp.gz initrd.vp  
rm -rf mnt  
cd /  
shutdown -r now
```

16. Once booted, verify that vpath devices are being used. Add all other paths and reboot again. The following commands can be used to verify the use of vpath devices.

```
mount  
swapon -s  
lsvpcfg  
datapath query device
```

Upgrading the SDD

Use the following procedure to upgrade SDD:

1. Place the SDD rpm update package on the system
2. Restart the system in single-path mode (no SDD)
3. You might come to a **Repair filesystem** prompt. If so, enter the root password. At the prompt, remount the file system as read/write.

```
mount -n -o remount,rw /
```

For systems that have `/boot` on a separate mount point, mount `/boot` partition using `/dev/sd` device.

4. Remove the previous SDD driver.

```
rpm -e IBMsdd
```

The `/etc/vpath.conf` file will be saved to `vpath.conf.rpmsave`.

5. Install the new SDD driver.

```

rpm -ivh IBMsdd-x.x.x.x-y.i686.rhel4.rpm
mkdir -p /boot/mnt
cd /boot
mv initrd.vp initrd.vp.gz
gunzip initrd.vp.gz
cd /boot/mnt
cpio -iv < ../initrd.vp
cp /opt/IBMsdd/sdd-mod.ko-`uname -r` /boot/mnt/lib/

```

6. Verify that the soft link `sdd-mod.ko` in `/boot/mnt/lib` points to the current SDD module.
7. Copy the new `cfgvpath` command and use the `ldd` command to verify that the correct libraries are installed for `/boot/mnt/opt/IBMsdd/bin/cfgvpath`.

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
```

8. Prepare the `initrd.vp` for booting.

```

cd /boot/mnt
find . | cpio -H newc -vo > /boot/initrd.vp
cd /boot
gzip initrd.vp
mv initrd.vp.gz initrd.vp
rm -rf mnt
cd /

```

9. Reboot the system.

SAN Boot instructions for RHEL 4 with IBM SDD (ppc)

The following procedure is used to install RHEL 4 ppc onto a System p host with fibre-channel connected storage. It is assumed that no installation is present to work from and when completed, the boot and swap devices will be running on IBM SDD vpath devices.

Prerequisite steps

1. Ensure that the following conditions exist before continuing with this procedure:
 - The installation target must be single-pathed before installing RHEL 4.
 - Have a copy of RHEL 4 either network accessible or on CD-ROM.
 - Be familiar with the RHEL 4 installation. This includes understanding which packages will be installed and how to select required options through the installation.
 - Be familiar with how to set up a SAN network or direct-attached SAN storage devices so that the host system can access LUNs from those storage systems (this procedure was performed on an ESS Model 800).
 - Be familiar with creating LUNs on the ESS Model 800 so that the host can access the ESS Model 800 devices. Although SDD functions correctly in single-path environments, there should be redundant physical paths to the devices from the host after installation of RHEL 4.
 - Optionally, have an understanding of how the Linux kernel boot process functions and what processes and procedures that are used to boot a Linux distribution for a local storage device.
 - Ensure that there is network access to the system.
2. Configure QLogic devices

- For ease of installation and to avoid issues with internal SCSI or IDE controllers, all internal disk drive controllers should be disabled. This procedure assumes that this has been done.
 - Verify that the QLogic SAN HBA devices that are configured for the host have been setup to have their BOOT BIOS enabled. This permits discovery and use of SAN disk devices during this procedure. While in the QLogic Utility, configure the ESS Model 800 device from which the system will boot. If the utility cannot see the correct device, check the SAN and ESS Model 800 configurations before continuing.
3. Configure root/boot/swap devices
 - The physical boot device that will be used for installation and booting should be at least 4 GB in size. This is the minimum size for installing all packages from the installation media to the boot devices.
 - It is also recommended that the swap device be at least the size of physical memory that is configured in the LPAR. For simplicity, these instructions assume that the root/boot/swap devices are all located on the same device; however, this is not a requirement for the installation.
 4. Installation Media

The installation media, or source for installation, can be CD-ROM, NFS, HTTP, FTP, or so on. For this installation, an NFS-exported set of CD-ROMs was used. You can use any of the installation sources that are listed.
 5. Use this procedure to install RHEL 4:
 - a. From the BIOS Menus, select the installation source to boot from.
 - b. Verify that the QLogic qla2030 SAN HBA module is loaded and that the SAN devices that will be used for installation have been detected successfully.

Note: Because of the way Linux discovers SAN devices, and if SAN devices have already been configured for multiple path access, Linux will discover the same physical device multiple times, once for each logical path to the device. Take note which device will be used for the installation before proceeding, that is, /dev/sda.
 - c. Select the desired options until arriving at the Installation Settings.

Here, modification of the partitioning settings is required for this installation. This is to make sure that the device noted in the previous step will be used for the root/boot installation target.
 - d. The details of installation and partitioning are not written up here. See the installation procedures to determine which packages are needed for the type of system being installed.
 6. Rebooting
 - a. On reboot, modify the BIOS to boot from hard disk, the system should now boot to the newly installed OS.
 - b. At this point the installed boot device can be set as the default boot device for the system. This step is not required, but is suggested because it enables unattended reboots after this procedure is complete.

SAN boot configuration

The system should now be booted in single-path mode from the SAN. Verify this by looking at the mounted devices and where the root is mounted, using the **df** command. Also verify that the swap, using **swapon -s** and other configured

partitions are correctly mounted. This completes the single-path boot from SAN. To modify the boot/root and other devices for booting using the SDD driver continue with the following steps.

The following list of suggestions should be noted before beginning this procedure:

Note:

1. All values and devices in the following procedure might not be the same on the system where this procedure is being conducted. It is, however, recommended (but not required) to use *vpatha* as the physical device for the root volume group.
2. Perform this procedure in a single-path environment. Once completed and booting with SDD configure the SAN for multipath.
3. All commands in this procedure begin with a # sign and might be followed by the output of that command.
4. This procedure uses a copy of the current initrd named *initrd.vp*.
5. Before starting SDD, comment out any sd devices from */etc/fstab* other than */boot*. This ensures that all devices are written to the */etc/vpath.conf* file. These devices can later be changed to vpath devices if the intent is to have them multi-pathed. This is not absolutely required.
6. The */etc/fstab* will also need to be modified to point */boot* from */dev/sd[x]* or *LABEL=[some_label_name_here]* to */dev/vpath[x]*.
7. Modify the */boot/grub/menu.lst* file to add an entry for the SDD initrd.
8. It is always a good idea to make copies of files that are going to be manually modified such as */etc/fstab*, */etc/vpath.conf* and */boot/grub/menu.lst*.
 1. Install the IBM SDD driver.

Download and install the IBM SDD driver for the kernel version being used. SDD is packaged in an RPM format and can be installed using the **rpm** command. See "Installing SDD" on page 205 for more information.
 2. Modify the */etc/fstab* file, ensuring that:
 - a. LABEL= is not being used
 - b. */boot* is mounted on a vpath deviceBecause Red Hat writes labels to the disk and uses labels in the */etc/fstab* the boot (*/boot*) device might be specified as a label, that is, LABEL=*/boot*. This might, however, be a different label other than LABEL=*/boot*. Check for line in the */etc/fstab* where */boot* is mounted and change it to the correct vpath device. Ensure that any other device specified with the LABEL= feature is changed to a */dev/sd* or */dev/vpath* device. LABEL= in a multi-pathed environment confuses Red Hat.

There is a one-to-one correlation between sd and vpath minor devices, such as, *sda1* and *vpatha1*. Major devices, however, might not necessarily correlate; for example, *sdb1* could be *vpathd1*. Because */boot* was installed on */dev/sda1* and *vpatha* corresponds to *sda* in the */etc/vpath.conf* file, */dev/vpatha1* will be the mount device for */boot*.
 3. Modify */etc/yaboot.conf*.

Add an entry for the SDD/LVM boot using *initrd.vp*.

```

image=/vmlinuz-2.6.9-22.0.1.EL
label=linux-22.01
read-only
initrd=/initrd-2.6.9-22.0.1.EL.img
append="console=hvc0 root=/dev/sda4"

image=/vmlinuz-2.6.9-22.0.1.EL
label=linux-22.01-sdd
read-only
initrd=/initrd.vp
append="console=hvc0 root=/dev/vpatha4"

```

Collect SDD data in preparation for configuring `/etc/vpath.conf`, `/etc/fstab/`, `etc/yaboot.conf` and `/boot/initrd`.

```
sdd start
```

The `/etc/vpath.conf` file has now been created. You must ensure that `vpatha` is the root device. Use the **`cfgvpath query device`** command to obtain the LUN ID of the root's physical device. (In this procedure, `sda` is the root device).

The **`cfgvpath query`** command produces output similar to the following example. Note that some data from the following output has been modified for ease of reading.

```

cfgvpath query

/dev/sda (8, 0) host=0 ch=0 id=0 lun=0 vid=IBM pid=2105800 serial=12020870
lun_id=12020870
/dev/sdb (8, 16) host=0 ch=0 id=0 lun=1 vid=IBM pid=2105800 serial=12120870
lun_id=12120870
/dev/sdc (8, 32) host=0 ch=0 id=0 lun=2 vid=IBM pid=2105800 serial=12220870
lun_id=12220870
/dev/sdd (8, 48) host=0 ch=0 id=0 lun=3 vid=IBM pid=2105800 serial=12320870
lun_id=12320870

```

The `lun_id` for `/dev/sdb` is 12020870. Edit the `/etc/vpath.conf` file using the `lun_id` for `vpatha`. Remove all other entries from this file (they will be automatically added later by SDD) .

4. Prepare the `initrd` file

The [initrd file] refers to the current `initrd` in `/boot`. The correct `initrd` can be determined by the following method:

```
ls -lA /boot | grep initrd | grep $(uname -r)
```

`initrd-2.6.9-22.0.1.img` is the result.

```

cd /boot
cp [initrd file] to initrd.vp.gz
gunzip initrd.vp.gz
mkdir /boot/mnt

```

5. Change directory to `/boot/mnt` and un-archive the `initrd` image to `/boot/mnt`.

```

cd /boot/mnt
cpio -iv < ../initrd.vp

```

6. Make additional directories in `/boot/mnt`.

```
mkdir mnt
mkdir -p opt/IBMsdd/bin
chmod -R 640 opt/IBMsdd
mkdir -p lib/tls
```

7. Copy the following `/etc` files to `/boot/mnt/etc`

```
cp /etc/vpath.conf /boot/mnt/etc/
cp /etc/group /boot/mnt/etc/
cp /etc/passwd /boot/mnt/etc/
cp /etc/nsswitch.conf /boot/mnt/etc/
cp /etc/fstab /boot/mnt/etc/
```

8. Copy `/opt/IBMsdd/bin/cfgvpath` to `/boot/mnt/opt/IBMsdd/bin/`

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
```

9. Copy required library files for `cfgvpath`. Use the `ldd` command to determine the library files and locations.

Example:

```
ldd /opt/IBMsdd/bin/cfgvpath | awk '{print $(NF-1)}'
```

The `ldd` command returns something similar to the following:

```
/lib/tls/libc.so.6
/lib/ld-linux.so.2
```

These file must be copied to the `/boot/mnt/lib64/tls/` and `/boot/mnt/lib64/` directories respectively.

10. Copy the correct `sdd-mod` to the `initrd` file system. Use the `uname -r` command to determine the correct `sdd-mod`.

The `uname -r` command returns something similar to `2.6.9-22.0.1`

```
cp /opt/IBMsdd/sdd-mod.ko-2.6.9-22.0.1 /boot/mnt/lib/sdd-mod.ko
```

11. Copy the following binaries and change permissions to `755`:

```
cp /bin/mknod /boot/mnt/bin/
cp /bin/mount /boot/mnt/bin/
cp /bin/umount /boot/mnt/bin/
cp /bin/cp /boot/mnt/bin/
chmod 755 /boot/mnt/bin/*
```

12. Copy required library files for each binary copied to the `/boot/mnt` directory in the previous step.

Use the `ldd` command to determine the library files and locations.

Many binaries use the same libraries so there might be duplications of copying.

Example:

```
ldd /bin/mknod | awk '{print $(NF-1)}' | grep lib

/lib/libselinux.so.1
/lib/tls/libc.so.6
/lib/ld-linux.so.2
```

The above files must be copied to the `/boot/mnt/lib/tls/` and `/boot/mnt/lib/` directories respectively.

Also, copy the following library files to `/boot/mnt/lib/`.

```
cp /lib/libproc-3.2.3.so /boot/mnt/lib/  
cp /lib/libtermcap.so.2 /boot/mnt/lib/  
cp /lib/libnss_files.so.2 /boot/mnt/lib/
```

13. Modify the `/boot/mnt/init` file.

Add the following lines after the modules load and just before the `/sbin/udevstart`. Note that `/sbin/udevstart` may exist multiple times in the `initrd`. Make sure these lines are added before the correct `/sbin/udevstart` entry which is located after the kernel modules load.

```
echo "Loading SDD module"  
insmod /lib/sdd-mod.ko  
echo "Creating vpath devices"  
/opt/IBMsdd/bin/cfgvpath
```

14. Use **cpio** to archive the `/boot/mnt` directory and **gzip** it in preparation for rebooting.

```
find . | cpio -H newc -vo > ../initrd.vp  
cd /boot  
gzip initrd.vp  
mv initrd.vp.gz initrd.vp  
rm -rf mnt  
cd /  
shutdown -r now
```

15. Install the yaboot boot loader to a bootstrap partition using the **ybin** command.

```
Ybin -b /dev/sda1
```

Where `/dev/sda1` is the PreP partition.

16. Verify System has reboot and SDD is configured correctly

Once booted, verify that vpath devices are being used. Add all other paths and reboot again. The following commands can be used to verify the use of vpath devices:

- `mount`
- `swapon -s`
- `lsvpcfg`
- `datapath query device`

Upgrading the SDD

Use the following procedure to upgrade SDD:

1. Place the SDD rpm update package on the system
2. Restart the system in single-path mode (no SDD)
3. You might come to a **Repair filesystem** prompt. If so, enter the root password. At the prompt, remount the file system as read/write.

```
mount -n -o remount,rw /
```

For systems that have `/boot` on a separate mount point, mount `/boot` partition using `/dev/sd` device.

4. Remove the previous SDD driver.

```
rpm -e IBMsdd
```

The `/etc/vpath.conf` file will be saved to `vpath.conf.rpmsave`.

5. Install the new SDD driver.

```
rpm -ivh IBMsdd-x.x.x.x-y.ppc64.rhel4.rpm
mkdir -p /boot/mnt
cd /boot
mv initrd.vp initrd.vp.gz
gunzip initrd.vp.gz
cd /boot/mnt
cpio -iv < ../initrd.vp
cp /opt/IBMsdd/sdd-mod.ko-`uname -r` /boot/mnt/lib/
```

6. Verify that the soft link `sdd-mod.ko` in `/boot/mnt/lib` points to the current `sdd` module.
7. Copy the new **cfgvpath** command and use the **ldd** command to verify that the correct libraries are installed for `/boot/mnt/opt/IBMsdd/bin/cfgvpath`.

```
cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
```

8. Prepare the `initrd.vp` for booting.

```
cd /boot/mnt
find . | cpio -H newc -vo > /boot/initrd.vp
cd /boot
gzip initrd.vp
mv initrd.vp.gz initrd.vp
rm -rf mnt
cd /
```

9. Install the `yaboot` boot loader to a bootstrap partition using the **ybin** command.

```
Ybin -b /dev/sda1
```

where `/dev/sda1` is the PreP partition.

10. Reboot the system.

SAN boot instructions for RHEL 4 with IBM SDD (x86) and LVM 2

Use this procedure to install RHEL 4 U1 (or later) and configure SDD with LVM. This procedure assumes that no installation is present to work from and when completed, the boot and swap devices will be running on SDD vpath devices and will be under LVM control.

Prerequisite steps

1. Ensure that the following conditions exist before continuing with this procedure:
 - The installation target **MUST** be single-pathed before installing RHEL 4.
 - A copy of RHEL 4 U1 i386 either network accessible or on CD.
 - Be familiar with the RHEL 4 installation. This includes understanding which packages will be installed.
 - Be familiar with setting up root and swap under LVM control.

- Be familiar with how to set up a SAN network or direct-attached SAN storage devices so that the host system can access LUNs from the those storage systems (this procedure was performed on an ESS Model 800).
 - Be familiar with creating LUNs on the ESS Model 800 so that the host can access the ESS Model 800 devices. Although SDD functions correctly in single path environments, it is recommended that there be redundant physical paths to the devices from the host after installation of RHEL 4.
 - Optionally, have an understanding of how the Linux kernel boot process functions and what processes and procedures that are used to boot a Linux distribution for a local storage device.
 - Ensure there will be network access to the system.
2. Configure the HBA devices.

Note: For ease of installation and to avoid issues with internal SCSI or IDE controllers, all internal disk drive controllers should be disabled. This procedure assumes that this has been done.

Verify that the SAN HBA devices that are configured for the host have been setup to have their BOOT BIOS enabled. This permits discovery and use of SAN disk devices during this procedure.

3. Configure the boot/root/swap devices.

The boot device that will be used for installation and booting should be at least 4 GB in size. This is the minimum size for installing a base package set from the installation media to the boot devices.

It is also recommended that the swap device be at least the size of physical memory that is configured in the host. For simplicity these instructions assume that the boot, root, and swap devices are all located on the same device; however, this is not a requirement for the installation.

The root (/) device must be under LVM control. The boot (/boot) device must NOT be under LVM control. Swap might also be under LVM control, but this is not a requirement. However, it should at least use a vpath device.

4. Use the installation media.

The installation media; that is, the source for installation, can be CD-ROM, NFS, HTTP, FTP, and so forth. For this installation, an NFS-exported set of CD-ROMs was used. You can use any of the installation sources that are listed.

5. Install the system.

- From the BIOS Menus select the installation source to boot from.
- Verify that the HBA module is loaded and that the SAN devices that will be used for installation have been detected successfully.

Note: Due to the way Linux discovers SAN devices, if SAN devices have already been configured for multiple path access, Linux will discover the same physical device multiple times, once for each logical path to the device. Note which device will be used for the installation before proceeding; that is, /dev/sda.

- Select the desired options until arriving at the Installation Settings. Here, modification of the partitioning settings is required for this installation. This is to make sure that the device noted in the previous step will be used for the root/boot installation target.

Note: The details of installation and partitioning are not documented here. See the installation procedures to determine which packages are needed for the type of system being installed.

6. Reboot the system.
 - a. On reboot, modify the BIOS to boot from hard disk. The system should now boot to the newly installed OS.
 - b. Verify that the system is booted from the correct disk and that the boot/root/swap and LVM configurations are correct.
 - c. At this point the installed boot device can be set as the default boot device for the system. This step is not required, but is suggested because it enables unattended reboots after this procedure is complete.

SAN boot configuration

The system should now be booted in single-path mode from the SAN. Verify this by looking at the mounted devices and where the root is mounted, using the **df** command. Also verify that the swap, using **swapon -s** and other configured partitions are correctly mounted. This completes the single-path boot from SAN.

Note:

1. All values and devices in the following procedure might not be the same on the system where this procedure is being conducted. Although it is not required, use vpatha as the physical device for the root volume group.
2. Perform this procedure in a single-path environment. Once completed and booting with SDD and LVM configure the SAN for multipath.
3. All commands in this procedure begin with a # sign and can be followed by the output of that command such as the command **pvdiskdisplay**.
4. Because /boot will not be under LVM control, it might be safer to work from within /boot.
5. This procedure uses a copy of the current initrd named initrd.vp.
6. The Volume groups for root and swap on this system are as follows:

```

/dev/rootVolGroup/
/dev/rootVolGroup/rootVol
/dev/rootVolGroup/swapVol
/dev/rootVolGroup/rootVol -> /dev/mapper/rootVolGroup-rootVol
/dev/rootVolGroup/swapVol -> /dev/mapper/rootVolGroup-swapVol
Physical device is sda2
vpath device vpatha2
```

7. Before starting SDD, comment out any SCSI disk devices from /etc/fstab other than /boot. This will ensure that all devices are written to the /etc/vpath.conf file. These devices can later be changed to vpath devices if the intent is to have them multipathed. This is not absolutely required.
8. The /etc/fstab will also need to be modified to point /boot from /dev/sd[x] or LABEL=[some_label_name_here] to /dev/vpath[x].
9. Modify the /boot/grub/menu.lst file to add an entry for the SDD initrd.
10. Modify /etc/lvm/lvm.conf to recognize vpath devices and ignore SCSI disk devices.
11. It is always a good idea to make copies of files that are going to be manually modified such as /etc/fstab, /etc/vpath.conf /etc/lvm/lvm.conf and /boot/grub/menu.lst.

To modify the boot/root and other devices for booting using the SDD driver continue with the following steps.

1. Install SDD driver IBMsdd-1.6.0.1-8.i686.rhel4.rpm
Change to the directory where the SDD rpm is located and use the rpm tool to install the IBMsdd driver and applications.

```
# rpm -ivh IBMsdd-1.6.0.1-8.i686.rhe14.rpm
```

2. Use **pvdisplay** to get the physical volume for the root and swap lvm volume group(s).

In this procedure /dev/sda2 (sda) is the device that will be used for /dev/vpatha2 (vpatha)

```
# pvdisplay
--- Physical volume ---
PV Name           /dev/sda2
VG Name           rootVolGroup
PV Size           9.09 GB / not usable 0
Allocatable       yes
PE Size (KByte)   32768
Total PE          291
Free PE           1
Allocated PE      290
PV UUID           SSm5g6-UoWj-evHE-kBj1-3QB4-EVi9-v88xiI
```

3. Modify the /etc/fstab file, ensuring that:

- a. LABEL= is not being used
- b. /boot is mounted on a vpath device

Because Red Hat writes labels to the disk and uses labels in the /etc/fstab, the boot (/boot) device might be specified as a label; that is, LABEL=/boot. This can, however, be a different label other than LABEL=/boot. Check for the line in the /etc/fstab where /boot is mounted and change it to the correct vpath device. Also ensure that any other device specified with the LABEL= feature is changed to a /dev/sd or /dev/vpath device. Red Hat does not recognize LABEL= in a multipathed environment.

There is a one-to-one correlation between SCSI disk and vpath minor devices, such as, sda1 and vpatha1. Major devices, however, might not correlate; that is, sdb1 could be vpathd1.

Because /boot was installed on /dev/sda1 and vpatha corresponds to sda in the /etc/vpath.conf file, /dev/vpatha1 will be the mount device for /boot.

Example:

Change:

```
/dev/rootVolGroup/rootVol /      ext3 defaults 1 1
LABEL=/boot                /boot ext3 defaults 1 2
/dev/rootVolGroup/swapVol swap swap defaults 0 0
```

To:

```
/dev/rootVolGroup/rootVol /      ext3 defaults 1 1
/dev/vpatha1                /boot ext3 defaults 1 2
/dev/rootVolGroup/swapVol swap swap defaults 0 0
```

4. Modify the /boot/grub/menu.lst file. Add an entry for the SDD/LVM boot using initrd.vp.

```

default=1
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Red Hat Enterprise Linux AS (2.6.9-11.ELsmp) w/LVM w/SDD
    root (hd0,0)
    kernel /vmlinuz-2.6.9-11.ELsmp ro root=/dev/rootVolGroup/rootVol
    initrd /initrd.vp
title Red Hat Enterprise Linux AS (2.6.9-11.ELsmp)
    root (hd0,0)
    kernel /vmlinuz-2.6.9-11.ELsmp ro root=/dev/rootVolGroup/rootVol
    initrd /initrd-2.6.9-11.ELsmp.img
title Red Hat Enterprise Linux AS-up (2.6.9-11.EL)
    root (hd0,0)
    kernel /vmlinuz-2.6.9-11.EL ro root=/dev/rootVolGroup/rootVol
    initrd /initrd-2.6.9-11.EL.img

```

5. Modify `/etc/lvm/lvm.conf`.

Change:

```
filter = [ "a.//" ]
```

To:

```
filter = [ "a/vpath*/", "r/sd*/" ]
```

In the `types =` section, add an entry for `vpath`

```
types = [ "vpath", 16 ]
```

6. Start SDD.

```
# sdd start
```

`/etc/vpath.conf` has now been created. You must ensure that `vpatha` is the root device. Use the `cfgvpath query device` command to obtain the LUN ID of the root's physical device. (In this procedure, `sda` is the root device).

The `cfgvpath query` command produces output similar to the following:

Note: Some data from the following output has been modified for ease of reading.

```

# cfgvpath query
/dev/sda (8, 0) host=0 ch=0 id=0 lun=0 vid=IBM pid=2105800 serial=12020870 lun_id=12020870
/dev/sdb (8, 16) host=0 ch=0 id=0 lun=1 vid=IBM pid=2105800 serial=12120870 lun_id=12120870
/dev/sdc (8, 32) host=0 ch=0 id=0 lun=2 vid=IBM pid=2105800 serial=12220870 lun_id=12220870
/dev/sdd (8, 48) host=0 ch=0 id=0 lun=3 vid=IBM pid=2105800 serial=12320870 lun_id=12320870

```

The `lun_id` for `/dev/sda` is 12020870. Edit the `/etc/vpath.conf` file using the `lun_id` for `vpatha`. Remove all other entries from this file. (They are automatically added later by SDD.)

```
vpatha 12020870
```

7. Prepare the `initrd` file.

The `[initrd file]` refers to the current `initrd` in `/boot`. The correct `initrd` can be determined by the following method:

```
# ls -lA /boot | grep initrd | grep $(uname -r)

# cd /boot
# cp [initrd file] to initrd.vp.gz
# gunzip initrd.vp.gz
# mkdir /boot/mnt
```

Note: The rest of this procedure involves working from /boot/mnt.

8. Change directory to /boot/mnt and unarchive the initrd image to /boot/mnt.

```
# cd /boot/mnt
# cpio -iv < ../initrd.vp
```

9. Make additional directories in /boot/mnt.

```
# mkdir mnt
# mkdir -p opt/IBMsdd/bin
# chmod -R 640 opt/IBMsdd
# mkdir -p lib/tls
```

10. Copy the following /etc files to /boot/mnt/etc.

```
# cp /etc/vpath.conf /boot/mnt/etc/
# cp /etc/group /boot/mnt/etc/
# cp /etc/passwd /boot/mnt/etc/
# cp /etc/nsswitch.conf /boot/mnt/etc/
# cp /etc/fstab /boot/mnt/etc/
# cp /etc/lvm/lvm.conf /boot/mnt/etc/lvm/
```

11. Modify the /boot/mnt/etc/nsswitch.conf file. (For rhel4u1i386, this might already be done.)

- a. Change:

```
passwd: compat
```

To:

```
passwd: files
```

- b. Change:

```
group: compat
```

To:

```
group: files
```

12. Copy /opt/IBMsdd/bin/cfgvpath to /boot/mnt/opt/IBMsdd/bin/ and change permissions to 755.

```
# cp /opt/IBMsdd/bin/cfgvpath /boot/mnt/opt/IBMsdd/bin/
# chmod 755 /boot/mnt/opt/IBMsdd/bin/*
```

13. Copy the required library files for **cfgvpath**. Use the **ldd** command to determine the library files and locations.

Example:

```
# ldd /opt/IBMsdd/bin/cfgvpath | awk '{print $(NF-1)}'
```

The **ldd** command returns the following:

```
/lib/tls/libc.so.6
/lib/ld-linux.so.2
```

For RHEL 4.0, copy these files to the `/boot/mnt/lib/tls/` and `/boot/mnt/lib/` directories respectively.

For RHEL 4.5, copy these files to the `/boot/mnt/lib64/tls/` and `/boot/mnt/lib64/` directories respectively.

- Copy the correct `sdd-mod` to the `initrd` file system. Use the **uname -r** command to determine the correct `sdd-mod`. **uname -r** returns `2.6.9-11.ELsmp`
For RHEL 4.0:

```
# cp /opt/IBMsdd/sdd-mod.ko-2.6.9-11.ELsmp /boot/mnt/lib/sdd-mod.ko
```

For RHEL 4.5:

```
# cp /opt/IBMsdd/sdd-mod.ko-2.6.9-11.ELsmp /boot/mnt/lib64/sdd-mod.ko
```

- Copy the following binaries and change permissions to `755`:

```
# cp /bin/mknod /boot/mnt/bin/
# cp /bin/mount /boot/mnt/bin/
# cp /bin/umount /boot/mnt/bin/
# cp /bin/cp /boot/mnt/bin/
# chmod 755 /boot/mnt/bin/*
```

- Copy the required library files for each binary. Use the **ldd** command to determine the library files and locations.

Note: Many binaries use the same libraries, so there might be duplications when copying.

Example:

```
# ldd /bin/mknod | awk '{print $(NF-1)}'

/lib/libselinux.so.1
/lib/tls/libc.so.6
/lib/ld-linux.so.2
```

For RHEL 4.0, copy these files to the `/boot/mnt/lib/tls/` and `/boot/mnt/lib/` directories respectively. Also, copy all of the following library files to `/boot/mnt/lib/`.

```
# cp /lib/libproc-3.2.3.so /boot/mnt/lib/
# cp /lib/libtermcap.so.2 /boot/mnt/lib/
# cp /lib/libnss_files.so.2 /boot/mnt/lib
```

For RHEL 4.5, copy these files to the `/boot/mnt/lib64/tls/` and `/boot/mnt/lib64/` directories respectively. Also, copy all of the following library files to `/boot/mnt/lib64/`.

```
# cp /lib/libproc-3.2.3.so /boot/mnt/lib64/
# cp /lib/libtermcap.so.2 /boot/mnt/lib64/
# cp /lib/libnss_files.so.2 /boot/mnt/lib64
```

- Modify the `/boot/mnt/init` file. Add the following lines just before the following statement:

```
[ echo "Loading dm-mod.ko module" ]
```

```
echo "Loading SDD module"  
insmod /lib/sdd-mod.ko  
echo "Creating vpath devices"  
/opt/IBMsdd/bin/cfgvpath
```

Ensure that an updated copy of vpath.conf is copied to the /root filesystem by using the following syntax to mount the root file system.

```
/bin/mount -o rw -t [fstype] [device] /mnt
```

Add the following lines just after [insmod /lib/dm-snapshot.ko]. The values used for the [fstype] and [device] here are only examples. Use the correct values for the system that is being configured.

```
/bin/mount -o rw -t ext3 /dev/rootVolGroup/rootVol /mnt  
/bin/cp /etc/vpath.conf /mnt/etc/  
/bin/umount /mnt
```

18. Use **cpio** to archive the /boot/mnt directory and **gzip** to compress it in preparation for rebooting.

```
# find . | cpio -H newc -vo > ../initrd.vp  
# cd /boot  
# gzip initrd.vp  
# mv initrd.vp.gz initrd.vp  
# cd /  
# shutdown -r now
```

19. Once booted, verify that vpath devices are being used. Add all other paths and reboot again. The following commands can be used to verify the use of vpath devices.

```
# mount  
# swapon -s  
# pvdisplay  
# lsvpcfg  
# datapath query device
```

Using lilo with SDD (remote boot) on x86

Use this section to boot SDD using the lilo boot loader in a remote/SAN boot environment.

This section only applies to booting using **lilo** and SDD vpath on x86-based systems. PowerPC-based systems use a combination of **lilo** and **yaboot**, whereas on x86-based systems, you rely solely on either **grub** or **lilo** as the boot loader.

When the **lilo** command runs, **lilo** reads the configuration file /etc/lilo.conf and attempts to read the disk geometry of the boot device in order to write its data to the master boot record. The method it uses for reading the disk geometry is dependent on a hardcoded list of major numbers that only supports certain types of devices. SDD vpath is not one of these supported devices. Therefore, when you are remote-booting and the MBR is contained on the vpath device, **lilo** is not able to complete. When running the command **lilo** to update the MBR, you will get this error message:

Fatal: Sorry, don't know how to handle device 0xMMmm, where *MM* is the major number and *mm* is the minor number of the device in question (in hex).

To prevent **lilo** from checking the major numbers, you can manually specify the geometry of the disk in the file `/etc/lilo.conf`.

Manually specifying disk geometry of the boot device

The syntax for manually specifying the disk geometry is:

```
disk = <disk-name>
  bios = 0x80
  sectors = <# of sectors/track>
  heads = <# of heads>
  cylinders = <# of cylinders>
  partition = <first vpath partition>
    start = <sector start # of first partition>
  partition = <second vpath partition>
    start = <sector start # of second partition>
  ...
  partition = <last vpath partition>
    start = <sector start # of last partition>
```

The following example is from a RHEL3 U4 system running on vpath devices:

```
disk = /dev/vpatha
  bios = 0x80
  sectors = 63
  heads = 255
  cylinders = 5221
  partition = /dev/vpatha1
    start = 63
  partition = /dev/vpatha2
    start = 1028160
  partition = /dev/vpatha3
    start = 79794855
```

Use the following procedure to find the information from your system in order to manually specify the disk geometry:

1. Use the **sfdisk** utility to find the cylinders, heads, and blocks. Use the `-l` option to list the current partition table and geometry numbers. For example,

```
[root@server ~]# sfdisk -l /dev/vpatha

Disk /dev/vpatha: 5221 cylinders, 255 heads, 63 sectors/track
Units = cylinders of 8225280 bytes, blocks of 1024 bytes, counting from 0

   Device Boot Start      End    #cyls   #blocks  Id System
/dev/vpatha1  *          0+        63        64-    514048+  83 Linux
/dev/vpatha2             64    4966    4903   39383347+  83 Linux
/dev/vpatha3          4967    5220     254    2040255    82 Linux swap
/dev/vpatha4             0         -         0         0      0 Empty
```

Note the cylinders, heads, and sectors per track and use this information to fill in the appropriate `lilo.conf` entries.

2. A separate program, **hdparm**, can be used to get the starting sector numbers for each partition.

However, **hdparm** only works on scsi disk or ide disk devices (`/dev/sdXX` or `/dev/hdXX`) and does not work on vpath devices. You can use one of the underlying paths that correspond to your boot disk to check the values. Your root disk is `vpatha` and there are four underlying SCSI disk devices or paths that correspond to that vpath device if your `lsvpcfg` output looks like the following example:

```
000 vpatha ( 254, 0) 75022811540 = 6005076303ffc06a0000000000001540 = /dev/sda /dev/sdb /dev/sdc /dev/sdd
```

3. Choose one vpath, for example, `/dev/sda`, and then issue the following command:

```
[root@server ~]# hdparm -g /dev/sda
```

The following output is displayed:

```
/dev/sda:  
geometry      = 5221/255/63, sectors = 83886080, start = 0
```

4. Compare this output to the `sfdisk -l` output.
5. Issue `hdparm -g` against every partition. For example:

```
[root@server ~]# hdparm -g /dev/sda  
  
/dev/sda:  
geometry      = 5221/255/63, sectors = 83886080, start = 0  
[root@server ~]# hdparm -g /dev/sda1  
  
/dev/sda1:  
geometry      = 5221/255/63, sectors = 1028097, start = 63  
[root@server ~]# hdparm -g /dev/sda2  
  
/dev/sda2:  
geometry      = 5221/255/63, sectors = 78766695, start = 1028160  
[root@server ~]# hdparm -g /dev/sda3  
  
/dev/sda3:  
geometry      = 5221/255/63, sectors = 4080510, start = 79794855
```

6. Use the values after the "start =" sections above as the starting sector numbers for the `/etc/lilo.conf` parameters. These values correspond to the starting sector numbers in the example snippet from `lilo.conf` shown in the example 290.
7. Insert the disk parameters and all the supporting information.
8. Rerun `lilo`.

The command should now succeed because it does not have to probe the geometry of the vpath device, but instead uses the entries in `lilo.conf`.

Here is an example of a `lilo.conf` file configured for remote boot:

```

boot=/dev/vpatha
map=/boot/map
install=/boot/boot.b
disk = /dev/vpatha
  bios = 0x80
  sectors = 63
  heads = 255
  cylinders = 5221
  partition = /dev/vpatha1
    start = 63
  partition = /dev/vpatha2
    start = 1028160
  partition = /dev/vpatha3
    start = 79794855
prompt
timeout=50
message=/boot/message
default=linux

image=/boot/vmlinuz-2.4.21-27.ELsmp
  label=linux
  initrd=/boot/initrd-2.4.21-27.ELsmp.img.test
  read-only
  root=/dev/vpatha2

```

SDD server daemon

The SDD server (also referred to as `sddsrv`) is an integrated component of SDD. This component consists of a UNIX application daemon that is installed in addition to the SDD device driver. See Chapter 9, “Using the SDDDSM or the SDDPCM daemon,” on page 365 for more information about `sddsrv`.

Verifying if the SDD server has started

After you have installed SDD, verify if the SDD server (`sddsrv`) has automatically started by entering `ps wax | grep sddsrv`.

If the SDD server (`sddsrv`) has automatically started, the output from the `ps` command looks like this:

```

31616 ? S 0:00 /opt/IBMsdd/bin/sddsrv
31617 ? S 0:00 /opt/IBMsdd/bin/sddsrv
31618 ? S 0:00 /opt/IBMsdd/bin/sddsrv
31619 ? S 0:10 /opt/IBMsdd/bin/sddsrv
31620 ? S 0:00 /opt/IBMsdd/bin/sddsrv
31621 ? S 0:00 /opt/IBMsdd/bin/sddsrv
31622 ? S 0:00 /opt/IBMsdd/bin/sddsrv

```

If processes are listed, the SDD server has automatically started.

If the SDD server has not started, no processes will be listed and you should see “Starting the SDD server manually” for instructions to start `sddsrv`.

Starting the SDD server manually

If the SDD server did not start automatically after you performed the SDD installation, use the following process to start `sddsrv`:

1. Edit `/etc/inittab` and append the following text:

```

#IBMsdd path recovery daemon:
srv:345:respawn:/opt/IBMsdd/bin/sddsrv > /dev/null 2>&1

```

2. Save the file `/etc/inittab`.
3. Enter the **telinit q** command.
4. Follow the directions in “Verifying if the SDD server has started” on page 292 to confirm that the SDD server started successfully.

Changing to a different port number for the SDD server

To change to a different port number for the SDD server, see the instructions in “Changing the `sddsrv` or `pcmsrv` TCP/IP port number” on page 369.

Stopping the SDD server

To stop the SDD server, complete the following steps:

1. Edit `/etc/inittab` and comment out the SDD server entry:

```
#IBMsdd path recovery daemon:  
#srv:345:respawn:/opt/IBMsdd/bin/sddsrv > /dev/null 2>&1
```

2. Save the file.
3. Issue **telinit q**.

See “Verifying if the SDD server has started” on page 292 to verify that the SDD server is not running. If `sddsrv` is not running, no processes will be listed when you enter **ps wax | grep sddsrv**.

Collecting trace information

SDD tracing can be enabled using the SDD server Web page. Enabling tracing puts the trace information into memory. To extract that information, issue **killall -IO sddsrv**. This command causes `sddsrv` to copy the trace data out of memory to the file `/var/log/sdd.log` on reception of this signal.

Understanding SDD support for single-path configuration

SDD does not support concurrent download of licensed machine code in single-path mode.

However, SDD supports *single-path* SCSI or fibre-channel connection from your Linux host system to a disk storage system and *single-path* fibre-channel connection from your Linux host system to a SAN Volume Controller.

Note:

1. SDD supports one fibre-channel adapter on the host system. SDD does not support SCSI adapters.
2. If your host has only one fibre-channel adapter port, it requires you to connect through a switch to multiple ports. You should have at least two fibre-channel adapters to prevent data loss due to adapter hardware failure or software failure for multipath support.
3. Because of single-path connection, SDD cannot provide single-point-failure protection and load balancing. IBM does not recommend this.

Partitioning SDD vpath devices

Disk partitions are known as *logical devices*. Disk partitions cannot be configured as SDD vpath devices; only entire SCSI disks can be configured. Once configured, an SDD vpath device can be partitioned into logical devices. The SDD naming scheme for disks and disk partitions follows the standard Linux disk-naming convention. The following description illustrates the naming scheme for SCSI disks and disk partitions:

1. The first two letters indicate the SCSI device.
2. The next letter (or two letters), a-z, specifies the unique device name.
3. A number following the device name denotes the partition number. For example, `/dev/sda` is the whole device, while `/dev/sda1` is a logical device representing the first partition of the whole device `/dev/sda`. Each device and partition has its own major and minor number.

Similarly then, a specific device file `/dev/vpathX` is created for each supported multipath SCSI disk device (where `X` represents the unique device name; as with SCSI disk devices, `X` might be one or two letters).

Device files `/dev/vpathXY` are also created for each partition of the multipath device (where `Y` represents the corresponding partition number). When a file system or user application wants to use the logical device, it should refer to `/dev/vpathXY` (for example, `/dev/vpatha1` or `/dev/vpathbc7`) as its multipath logical device. All I/O management, statistics, and failover processes of the logical device follow those of the whole device.

The following output demonstrates how the partitions are named:

```
brw-r--r-- 1 root  root  247,  0 Apr 2 16:57 /dev/vpatha
brw-r--r-- 1 root  root  247,  1 Apr 2 16:57 /dev/vpatha1
brw-r--r-- 1 root  root  247, 10 Apr 2 16:57 /dev/vpatha10
brw-r--r-- 1 root  root  247, 11 Apr 2 16:57 /dev/vpatha11
brw-r--r-- 1 root  root  247, 12 Apr 2 16:57 /dev/vpatha12
brw-r--r-- 1 root  root  247, 13 Apr 2 16:57 /dev/vpatha13
brw-r--r-- 1 root  root  247, 14 Apr 2 16:57 /dev/vpatha14
brw-r--r-- 1 root  root  247, 15 Apr 2 16:57 /dev/vpatha15
brw-r--r-- 1 root  root  247,  2 Apr 2 16:57 /dev/vpatha2
brw-r--r-- 1 root  root  247,  3 Apr 2 16:57 /dev/vpatha3
brw-r--r-- 1 root  root  247,  4 Apr 2 16:57 /dev/vpatha4
brw-r--r-- 1 root  root  247,  5 Apr 2 16:57 /dev/vpatha5
brw-r--r-- 1 root  root  247,  6 Apr 2 16:57 /dev/vpatha6
brw-r--r-- 1 root  root  247,  7 Apr 2 16:57 /dev/vpatha7
brw-r--r-- 1 root  root  247,  8 Apr 2 16:57 /dev/vpatha8
brw-r--r-- 1 root  root  247,  9 Apr 2 16:57 /dev/vpatha9
```

Note: For supported file systems, use the standard UNIX `fdisk` command to partition SDD vpath devices.

Using standard UNIX applications

After successful installation, SDD resides above the SCSI subsystem in the block I/O stack of the Linux host system. In other words, SDD recognizes and communicates with the native device driver of your Linux host system and standard UNIX applications, such as `fdisk`, `fsck`, `mkfs`, and `mount` accept an SDD device name as a parameter. Therefore, SDD vpath device names can replace corresponding SCSI disk device name entries in system configurations files, such as `/etc/fstab`.

Make sure that the SDD devices match the devices that are being replaced. You can issue the `lsvpcfg` command to list all SDD devices and their underlying disks.

Managing common issues

The following list includes known general issues.

- Understanding SDD error recovery policy during total path failure
To avoid potential data corruption, if a write I/O fails on all paths of an SDD vpath device, all the paths are put into DEAD state and OFFLINE mode. The only way to make paths available again is through manual intervention. That is, you have to unmount the SDD vpath device and, if necessary, run `fsck` to check and repair the filesystem if necessary.
- A DEAD path is not reclaimed by sddsrvc because it was offlined by the kernel.
Sometimes, such as when a failure occurs or when there is a lot of I/O, the kernel's SCSI layer will offline a path. This behavior is done independently of SDD. Therefore, sddsrvc cannot reclaim a path that has been taken offline by the SCSI layer. Typically the system messages log contain messages indicating that this has happened. See Appendix B, "System log messages," on page 397 for more information about the messages that are logged for AIX SDD.

For example:

On Linux 2.4 kernels:

```
kernel: scsi: device set offline - not ready or  
command retry failed after bus reset: host 2 channel 0 id 1 lun 0
```

On Linux 2.6 kernels:

```
kernel: scsi: Device offlined - not ready after error recovery: host 2 channel 0 id 1 lun 0
```

or,

```
scsil (1:1): rejecting I/O to offline device
```

For 2.4 kernels, the only way to restore devices that have been taken offline by the SCSI layer is to reload the HBA driver.

For 2.6 kernels, you can use the sysfs interface to dynamically re-enable SCSI devices that have been taken offline by the SCSI layer.

- Setting SCSI midlayer timeout values to address loaded storage targets
Some storage devices require a longer time period to retire an I/O command issued by an initiator under heavy load. By default, the SCSI midlayer allots only 30 seconds per SCSI command before cancelling the I/O command to the initiator. Consider setting the timeout value to 60 seconds.

If you see SCSI errors of value 0x6000000, LUN reset messages or abort I/O messages, you can change the timeout setting in case it helps to alleviate the situation. It might also be necessary to stop all I/O operations and allow the target to retire all outstanding I/O before starting I/O operations again with the new timeout.

For Linux 2.6 kernels, you can manually set the timeout value through the sysfs interface. To do this, issue the following command: `echo 60 >`

`/sys/class/scsi_device/<host>:<channel>:<target>:<lun>/timeout`

where, replace the items in <> with the following (you can match with the values in `/proc/scsi/scsi`):

– *host* - host number

- *channel* - channel number
- *target* - target number
- *lun* - lun number

To simplify this process for multiple paths, Emulex has provided the script **set_timeout_target.sh** at the Emulex website under the Linux tools page. Because this script deals with SCSI disk devices, it can work equally well in environments that use Qlogic host bus adapters. Details on how to use the tool are available on the Emulex website: www.emulex.com/files/downloads/linux/tools.html#.

- Changing default queue depth values to avoid overloaded storage targets
You should lower the queue depth per LUN when using multipathing. With multipathing, this default value is magnified because it equals the default queue depth of the adapter multiplied by the number of active paths to the storage device. For example, given that Qlogic uses a default queue depth of 32, the recommended queue depth value to use would be 16 when using two active paths and 8 when using four active paths. Directions for adjusting the queue depth is specific to each HBA driver and should be available in the documentation for the HBA.

Chapter 6. Using the SDD on a NetWare host system

This chapter provides step-by-step procedures on how to install, configure, upgrade, and remove SDD on a NetWare host system (NetWare 5.1, NetWare 6.0, or NetWare 6.5) that is attached to a disk storage system. The SDD for NetWare is shipped as a Novell Custom Device Module (CDM), which is a driver component that is associated with storage devices and the commands that control the storage device. For updated and additional information not included in this chapter, see the Readme file on the CD-ROM or visit the following SDD website:

www-01.ibm.com/support/docview.wss?uid=ssg1S4000118

Attention: SDD does not support Novell NetWare host systems attached to SAN Volume Controller.

Verifying the hardware and software requirements

You must have the following hardware and software components in order to successfully install SDD.

You can check for and download the latest APARs, maintenance level fixes, and microcode updates from the following website:

www-933.ibm.com/support/fixcentral/

Hardware requirements

The following hardware components are needed:

- IBM TotalStorage SAN Fibre Channel Switch 2109 is recommended
- Host system
- Fibre-channel switch
- SCSI adapters and cables (ESS)
- Fibre-channel adapters and cables

Software requirements

The following software components are needed:

- Microsoft Windows operating system running on the client
- One of the following NetWare operating systems running on the server:
 - Novell NetWare 5.1 with Support Pack
 - Novell NetWare 6 with Support Pack
 - NetWare 6.5 with Support Pack
- NetWare Cluster Service for NetWare 5.1 if servers are being clustered
- NetWare Cluster Service for NetWare 6.0 if servers are being clustered
- NetWare Cluster Service for NetWare 6.5 if servers are being clustered
- ConsoleOne
- SCSI and fibre-channel device drivers

Supported environments

SDD supports the following environments:

- Novell NetWare 5.1 SP6
- Novell NetWare 6 SP1, SP2, SP3, SP4, or SP5
- Novell NetWare 6.5 SP1.1 or SP2
- Novell Cluster Services 1.01 for Novell NetWare 5.1 is supported on fibre-channel and SCSI devices.
- Novell Cluster Services 1.6 for Novell NetWare 6.0 is supported only for fibre-channel devices.
- Novell Cluster Services 1.7 for Novell NetWare 6.5 is supported only for fibre-channel devices.

Currently, only the following QLogic fibre-channel adapters are supported with SDD:

- QL2310FL
- QL2200F
- QLA2340 and QLA2340/2

Unsupported environments

SDD does not support:

- A host system with both a SCSI and fibre-channel connection to a shared disk storage system LUN
- Single-path mode during concurrent download of licensed machine code nor during any disk storage system concurrent maintenance that impacts the path attachment, such as a disk storage system host-bay-adapter replacement
- DS8000 and DS6000 do not support SCSI connectivity.

Disk storage system requirements

To successfully install SDD, ensure that the disk storage system devices are configured as one of the following:

- IBM 2105xxx (SCSI-attached device) or IBM FC 2105 (fibre-channel-attached device), for ESS devices
- IBM 2107xxx, for DS8000 devices
- IBM 1750xxx, for DS6000 devices
- IBM 2145, for SAN Volume Controller devices

where xxx represents the disk storage system model number.

SCSI requirements

To use the SDD SCSI support, ensure that your host system meets the following requirements:

- A SCSI cable connects each SCSI host adapter to an ESS port.
- If you need the SDD I/O load-balancing and failover features, ensure that a minimum of two SCSI adapters are installed.

For information about the SCSI adapters that can attach to your NetWare host system, go to the following website:

www.ibm.com/servers/storage/support

Fibre-channel requirements

You must check for and download the latest fibre-channel device driver APARs, maintenance level fixes, and microcode updates from the following website:

www-933.ibm.com/support/fixcentral/

Note: If your host has only one fibre-channel adapter, you need to connect through a switch to multiple disk storage system ports. You should have at least two fibre-channel adapters to prevent data loss due to adapter hardware failure or software failure.

To use the SDD fibre-channel support, ensure that your host system meets the following requirements:

- The NetWare host system has the fibre-channel device drivers installed.
- A fiber-optic cable connects each fibre-channel adapter to a disk storage system port.
- If you need the SDD I/O load-balancing and failover features, ensure that a minimum of two paths to a device are attached.

For information about the fibre-channel adapters that can be used on your NetWare host system, go to the following website:

www.ibm.com/servers/storage/support

Preparing for SDD installation

Before you install SDD, you must configure the disk storage system for your host system and attach required fibre-channel adapters.

Configuring the disk storage system

Before you install SDD, you must configure the disk storage system:

- The disk storage system to your host system and the required fibre-channel that are attached.
- The ESS to your host system and the required SCSI adapters that are attached.
- The disk storage system for single-port or multiple-port access for each LUN.

SDD requires a minimum of two independent paths that share the same logical unit to use the load-balancing and path-failover-protection features. **With a single path, failover protection is not provided.**

See the Installation and Planning Guide for your disk storage system for more information about how to configure the disk storage system.

See the Host Systems Attachment Guide for your disk storage system for information on working around Novell LUN limitations.

Configuring fibre-channel adapters

You must configure the fibre-channel adapters and the adapters' drivers that are attached to your NetWare host system before you install SDD. Follow the adapter-specific configuration instructions to configure the adapters.

For QLogic adapters, you need to add `/LUNS`, `/ALLPATHS`, `/PORTNAMES` while loading FC HBA device driver. For example:

```
LOAD QL2200.HAM SLOT=x /LUNS /ALLPATHS /PORTNAMES /GNNFT
LOAD QL2200.HAM SLOT=y /LUNS /ALLPATHS /PORTNAMES /GNNFT
```

Modify the startup.ncf file by adding SET MULTI-PATH SUPPORT=OFF at the top.

Then, modify the autoexec.ncf by adding SCAN ALL LUNS before MOUNT ALL:

```
...
...
SCAN ALL LUNS
MOUNT ALL
...
...
```

Ensure that you can see all the LUNs before installing SDD. Use the **list storage adapters** command to verify that all the LUNs are available.

See the *IBM TotalStorage Enterprise Storage Server Host Systems Attachment Guide* for more information about how to install and configure fibre-channel adapters for your NetWare host system.

See the *IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide* for working around NetWare LUN limitations.

Configuring SCSI adapters

Before you install and use SDD, you must configure your SCSI adapters.

For Adaptec AHA2944 adapters, add LUN_ENABLE=FFFF in startup.ncf:

```
LOAD AHA2940.HAM slot=x LUN_ENABLE=FFFF
LOAD AHA2940.HAM slot=y LUN_ENABLE=FFFF
```

See the *IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide* for more information about how to install and configure fibre-channel adapters for your NetWare host system.

See the *IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide* for information about working around NetWare LUN limitations.

Using a NetWare Compaq Server

When SDD is installed on a Compaq server running Novell NetWare, SDD may not failover as designed. Volume dismounts, hangs, or abnormal ends can result. Compaq servers running Novell NetWare can be configured to load the Compaq-specific CPQSHD.CDM driver. This driver has different behavior than the standard Novell SCSIHD.CDM driver. The CPQSHD.CDM driver will often do a re-scan after a path is lost. This re-scan can potentially cause volumes to be dismounted, and hangs or abnormal ends can result.

To ensure that SDD failover functions as designed and to prevent potential volume dismounts, hangs, or abnormal ends, do not load the CPQSHD.CDM file at startup. Remove the reference to this file from the STARTUP.NCF file or by comment out the line which loads CPQSHD.CDM. The standard Novell SCSIHD.CDM driver must be loaded in the place of the Compaq CPQSHD.CDM file startup. For example, the STARTUP.NCF file should look similar to the following example in order for SDD to failover as designed on a Novell NetWare Compaq server:

```
SET MULTI-PATH SUPPORT = OFF
...
#LOAD CPQSHD.CDM
...
LOAD SCSIHD.CDM
...
LOAD QL2300.HAM SLOT=6 /LUNS /ALLPATHS /PORTNAMES /GNNFT
LOAD QL2300.HAM SLOT=5 /LUNS /ALLPATHS /PORTNAMES /GNNFT
```

Using SCSIHD.CDM rather than CPQSHD.CDM will not cause any problems when running SDD on a Novell NetWare Compaq server.

Installing SDD

This section describes how to install SDD from CD-ROM and downloaded code.

Installing SDD from CD-ROM

The installation CD-ROM contains the following files:

- INSTALL.NLM, main body of the installer that contains the startup program
- SDD.CDM, a device driver
- DATAPATH.NLM, datapath command
- COPY.INS, the file copy destination
- STARTUP.INS, the STARTUP update
- INFO.INS, contains messages displayed at installation
- AUTOEXEC.INS, unused

To install the SDD:

1. Insert the SDD installation media into the CD-ROM drive.
2. Enter **load XXX:\path \install**, where XXX is the name of the CD volume mounted, in the NetWare console window to invoke INSTALL.NLM. This file starts the installation, copies SDD.CDM to a target directory, and updates the startup file.

Installing SDD from downloaded code

Use this procedure to install SDD from downloaded code:

1. Download the NetWare SDD package from the web as a zipped file.
2. Unzip the package and copy the files to the desired location.
3. Enter **load XXX:\path \install**, where XXX is the location for the SDD installation files, in the NetWare console window to invoke INSTALL.NLM. This file starts the installation, copies SDD.CDM to a target directory, and updates the startup file.

Configuring SDD

To load the SDD module, enter **load SDD**.

To unload the SDD module, enter **unload SDD**.

Note: During the initial load of SDD, NetWare SDD takes over control of each LUN (path) from other device drivers in order to virtualize all the underlying paths (LUNs). Messages like the following are issued during the initial load:

Device " {V597-A2-D1:2} IBM 2105F20 REV:.105" deactivated by driver due to device failure.

These messages are expected and are not cause for concern.

Maximum number of LUNs

SDD supports a total of less than 600 devices. The total devices supported equals the number of LUNs multiplied by the number of paths per LUN.

Displaying the current version of the SDD

Enter **modules SDD** to display the current version of the SDD.

Features

SDD provides the following functions:

- Automatic path detection, failover and selection
- Manual operations (datapath command)
- Path selection algorithms
- Dynamic load balancing
- Disk storage system logical unit detection
- Error reporting and logging
- SDD in NetWare-layered architecture

Automatic path detection, failover and selection

The SDD failover-protection system is designed to minimize any disruptions in I/O operations from a failing datapath. When a path failure is detected, the SDD moves the I/O access to another available path in order to keep the data flow. The SDD has the following path states:

- OPEN state
- CLOSE (Error) state
- DEAD state
- INVALID (PERMANENTLY DEAD) state

The OPEN state indicates that a path is available. This is the initial path state after the system starts. When a path failure occurs in the OPEN state, the path is put into the CLOSE (Error) state. If the SDD recovers the path, the path is put back into the OPEN state. While path recovery is in progress, the path is temporarily changed to the OPEN state.

If a path failure occurs three consecutive times in the CLOSE (Error) state, the path is put into the DEAD state in multipath mode. In the single-path mode, it stays in the CLOSE state. However, if the path is recovered, it is put back into the OPEN state. While path reclamation is in progress, the path is temporarily changed to OPEN state. The path is put into the INVALID state and is placed offline if path reclamation fails.

Only a datapath command, **datapath set adapter <n> online** or **datapath set device <n> path <m> online**, can return the path to the OPEN state.

In the event that all the paths fail, all the paths except one are moved into the DEAD state. The one path will still be in OPEN state. This indicates that further access to

LUNs is still accepted. At each access, all paths are attempted until at least one of them is recovered. The error count is incremented only for the path in the OPEN state while all other paths are failed.

Manual operations using the datapath commands

The **datapath** commands allow manual path selection using a command line interface. See Chapter 10, “Using the datapath commands,” on page 371 for detailed information about the commands.

SDD in the Novell NetWare environment supports the **datapath set device policy** command, which has the following options:

- **rr**, where *rr* indicates round robin
- **lb**, where *lb* indicates load balancing
- **df**, where *df* indicates the default policy, which is round robin
- **fo**, where *fo* indicates failover policy

Note: The **rr**, **lb**, and **df** options currently have the same effect.

The path-selection policy algorithms are:

round robin (rr)

The path to use for each I/O operation is chosen at random from paths that were not used for the last I/O operation. If a device has only two paths, SDD alternates between the two.

load balancing (lb)

The path to use for an I/O operation is chosen by estimating the load on the adapter to which each path is attached. The load is a function of the number of I/O operations currently in process. If multiple paths have the same load, a path is chosen at random from those paths.

default

This is the round-robin-path operation and is the default value.

failover only (fo)

All I/O operations for the device are sent to the same (preferred) path until the path fails because of I/O errors. Then an alternate path is chosen for subsequent I/O operations.

The **datapath open device path** command, which is supported on other platforms, is not supported in NetWare because it is not possible to open a device that failed to configure in NetWare.

NetWare does support the **scan** command, which scans the devices connected to the server. In case a device is detected, a message is sent to the SDD, and the SDD updates the path configuration based on the message. Therefore, you should issue the **scan all** command manually instead of the **addpath** command used on other platforms. You can also use the **scan all** command to put a new path under SDD control. **scan all** refreshes the device table and sends a message to the SDD in case a new device is found. SDD checks to see if the new device is a LUN under the disk storage system and, if so, adds it to the path group.

See Chapter 10, “Using the datapath commands,” on page 371

for more information about the **datapath** commands.

Understanding SDD error recovery algorithms

SDD assumes the following two operation modes:

- Single-path mode
- Multiple-path mode

Single-path mode

In single-path mode, only a single path is available in access to a device in a subsystem. The SDD never puts this path into the DEAD state.

Multiple-path mode

In this mode, two or more paths are available in access to a device in a subsystem. SDD has the following behavior concerning path operations:

- After a path failure occurs on a path, SDD attempts to use the path again after 2 000 successful I/O operations through another operational path or paths. This process is called *Path Recovery*.
- If the consecutive error count on the path reaches three, SDD puts the path into the DEAD state.
- SDD reverts the failed path from the DEAD state to the OPEN state after 50 000 successful I/O operations through an operational path or paths. This process is called *Path Reclamation*.
- If an access fails through the path that has been returned to the OPEN state, SDD puts the path into the INVALID (PERMANENTLY DEAD) state and then never attempts the path again. Only a manual operation using a **datapath** command can reset a path from the PERMANENTLY DEAD state to the OPEN state.
- All knowledge of prior path failures is reset when a path returns to the OPEN state.
- SDD never puts the last operational path into the DEAD state. If the last operational path fails, SDD attempts a previously-failed path or paths even though that path (or paths) is in PERMANENTLY DEAD state.
- If all the available paths failed, SDD reports an I/O error to the application.
- If the path is recovered as either a path recovery operation or a path reclamation operation, the path is then handled as a normal path in the OPEN state and the SDD stops keeping a history of the failed path.

Note: You can display the error count with the **datapath** command.

Dynamic load balancing

SDD distributes the I/O accesses over multiple active paths, eliminating data path bottlenecks.

Disk storage system logical unit detection

SDD works only with disk storage system logical units. SDD assumes that all logical units have 2105 as their first four characters in the Product ID in Inquiry Data. The Product ID indicates that it is a logical unit. The SDD also assumes that all logical units return unique serial numbers regardless of a port on the disk storage system.

Error reporting and logging

All error reports generated by SDD are logged in a NetWare standard log file, SYS:\SYSTEM\SYS\$LOG.ERR. Any path state transition is logged in the log file. The log has the following information:

- Event source name
- Time stamp
- Severity
- Event number
- Event description
- SCSI sense data (in case it is valid)

Note: A failure in Path Recovery or Path Reclamation is not logged, while a successful path recovery in Path Recovery or Path Reclamation is logged.

SDD in NetWare-layered architecture

All path-management features are implemented in an SDD-unique Custom Device Module (CDM), which is called SDD.CDM. It supports LUNs under disk storage systems only. Any other LUNs are supported by a NetWare standard CDM, SCSIHD.CDM. The SDD.CDM has all functions that the standard CDM has in addition to the disk storage system-specific path management features. The SDD.CDM assumes that it will be working with a standard Host Adapter Module (HAM).

NetWare has assigned the SDD CDM module ID 0x7B0.

Display a single device for a multipath device

With SDD version 1.00i, the system will display a single device for a multipath device. However, **datapath query device** will show all the paths for each device. For example, with older versions of SDD, on a system with two LUNs with each having two paths, the following output would be displayed for the **list storage adapters** command:

```
V597-A3] QL2300 PCI FC-AL Host Adapter Module
[V597-A3-D0:0] IBM 2105800 rev:.324
[V597-A3-D0:1] IBM 2105800 rev:.324
[V597-A4] QL2300 PCI FC-AL Host Adapter Module
[V597-A4-D0:0] IBM 2105800 rev:.324
[V597-A4-D0:1] IBM 2105800 rev:.324
```

Starting with SDD version 1.00i, the **list storage adapters** displays:

```
[V597-A3] QL2300 PCI FC-AL Host Adapter Module
[V597-A3-D0:0] IBM 2105800 rev:.324
[V597-A3-D0:1] IBM 2105800 rev:.324
[V597-A4] QL2300 PCI FC-AL Host Adapter Module
```

The **datapath query device** output will be same in both the cases.

Removing the SDD

Complete the following steps to remove the SDD:

1. Manually remove files from the C:\NWSERVER directory.
2. Remove SDD-related entries in startup.ncf.

Cluster setup for Novell NetWare 5.1

To set up clustering in Novell NetWare 5.1, follow the steps described in the Novell Cluster Services document available online at: www.novell.com/documentation/lg/ncs/index.html

Cluster setup for Novell NetWare 6.0

To set up clustering in Novell NetWare 6.0, follow the steps described in the Novell Cluster Services document available online at: www.novell.com/documentation/lg/ncs6p/index.html

Examples of commands output on the Console Window

The following examples show the basic commands output during path failover and failback. The examples are from NetWare 6.0 SP2.

```
END:modules sdd
SDD.CDM
Loaded from [C:\NWSERVER\]
(Address Space = OS)
IBM Enterprise Storage Server SDD CDM
Version 1.00.07 July 17, 2003
(C) Copyright IBM Corp. 2002 Licensed Materials - Property of IBM
END:datapath query device
Total Devices : 2
DEV#: 3A DEVICE NAME: 0x003A:[V596-A4-D1:0]
TYPE: 2105E20 SERIAL: 30812028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003A:[V596-A4-D1:0] OPEN NORMAL 14 0
1 0x007A:[V596-A3-D1:0] OPEN NORMAL 14 0
2 0x001A:[V596-A4-D0:0] OPEN NORMAL 14 0
3 0x005A:[V596-A3-D0:0] OPEN NORMAL 14 0
DEV#: 3B DEVICE NAME: 0x003B:[V596-A4-D1:1]
TYPE: 2105E20 SERIAL: 01312028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003B:[V596-A4-D1:1] OPEN NORMAL 1 0
1 0x007B:[V596-A3-D1:1] OPEN NORMAL 1 0
2 0x001B:[V596-A4-D0:1] OPEN NORMAL 1 0
3 0x005B:[V596-A3-D0:1] OPEN NORMAL 0 0
END:datapath query adapter
Active Adapters :2
Adpt# Adapter Name State Mode Select Errors Paths Active
0 [V596-A4] NORMAL ACTIVE 30 0 4 4
1 [V596-A3] NORMAL ACTIVE 30 0 4 4
(Creating volume tempvol on DEV#3A through ConsoleOne, mount tempvol)
END:mount tempvol
Activating volume "TEMPVOL"
** Volume layout v35.00
** Volume creation layout v35.00
** Processing volume purge log
** .
Volume TEMPVOL set to the ACTIVATE state.
Mounting Volume TEMPVOL
** TEMPVOL mounted successfully
END:volumes
Mounted Volumes Name Spaces Flags
SYS DOS, LONG Cp Sa
_ADMIN DOS, MAC, NFS, LONG NSS P
TEMPVOL DOS, MAC, NFS, LONG NSS
3 volumes mounted
(start IO)
END:datapath query device
```

```

Total Devices : 2
DEV#: 3A DEVICE NAME: 0x003A:[V596-A4-D1:0]
TYPE: 2105E20 SERIAL: 30812028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003A:[V596-A4-D1:0] OPEN NORMAL 224 0
1 0x007A:[V596-A3-D1:0] OPEN NORMAL 224 0
2 0x001A:[V596-A4-D0:0] OPEN NORMAL 224 0
3 0x005A:[V596-A3-D0:0] OPEN NORMAL 224 0
DEV#: 3B DEVICE NAME: 0x003B:[V596-A4-D1:1]
TYPE: 2105E20 SERIAL: 01312028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003B:[V596-A4-D1:1] OPEN NORMAL 1 0
1 0x007B:[V596-A3-D1:1] OPEN NORMAL 1 0
2 0x001B:[V596-A4-D0:1] OPEN NORMAL 1 0
3 0x005B:[V596-A3-D0:1] OPEN NORMAL 1 0
END:datapath query adapter
Active Adapters :2
Adpt# Adapter Name State Mode Select Errors Paths Active
0 [V596-A4] NORMAL ACTIVE 795 0 4 4
1 [V596-A3] NORMAL ACTIVE 794 0 4 4
(Pull one of the cable)
Error has occurred on device 0x3A path 2
(Adapter Error Code: 0x8007, Device Error Code: 0x0000)
[No sense data]
This path is in CLOSE state.
Error has occurred on device 0x3A path 0
(Adapter Error Code: 0x8007, Device Error Code: 0x0000)
[No sense data]
This path is in CLOSE state.
Path Recovery (1) has failed on device 0x3A path 2
(Adapter Error Code: 0x8007, Device Error Code: 0x0000)
[No sense data]
This path is in CLOSE state.
Path Recovery (1) has failed on device 0x3A path 0
(Adapter Error Code: 0x8007, Device Error Code: 0x0000)
[No sense data]
This path is in CLOSE state.
ND:datapath query device
Total Devices : 2
DEV#: 3A DEVICE NAME: 0x003A:[V596-A4-D1:0]
TYPE: 2105E20 SERIAL: 30812028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003A:[V596-A4-D1:0] CLOSE NORMAL 418 2
1 0x007A:[V596-A3-D1:0] OPEN NORMAL 740 0
2 0x001A:[V596-A4-D0:0] CLOSE NORMAL 418 2
3 0x005A:[V596-A3-D0:0] OPEN NORMAL 739 0
DEV#: 3B DEVICE NAME: 0x003B:[V596-A4-D1:1]
TYPE: 2105E20 SERIAL: 01312028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003B:[V596-A4-D1:1] OPEN NORMAL 1 0
1 0x007B:[V596-A3-D1:1] OPEN NORMAL 1 0
2 0x001B:[V596-A4-D0:1] OPEN NORMAL 1 0
3 0x005B:[V596-A3-D0:1] OPEN NORMAL 1 0
END:datapath query adapter
Active Adapters :2
Adpt# Adapter Name State Mode Select Errors Paths Active
0 [V596-A4] DEGRAD ACTIVE 901 5 4 2
1 [V596-A3] NORMAL ACTIVE 1510 0 4 4
(If reconnect cable and issue manual online command)
END:datapath set adapter 0 online
datapath set adapter command has been issued for adapter 4(Adpt# 0).
This adapter is in NORMAL state.
device 0x59 path 0 is in OPEN state.
device 0x58 path 0 is in OPEN state.
datapath set adapter command has been issued for adapter 4(Adpt# 2).
This adapter is in NORMAL state.

```

```

device 0x59 path 2 is in OPEN state.
device 0x58 path 2 is in OPEN state.
Success: set adapter 0 to online
Adpt# Adapter Name State Mode Select Errors Paths Active
0 [V596-A4] NORMAL ACTIVE 2838 14 4 4
(If reconnect cable and let SDD do path recovery itself)
Path Recovery (2) has succeeded on device 0x3A path 2.
This path is in OPEN state.
Path Recovery (2) has succeeded on device 0x3A path 0.
This path is in OPEN state.
(If cable is not reconnected, after 3 retries, path will be set to DEAD)
Path Recovery (3) has failed on device 0x3A path 2
(Adapter Error Code: 0x8007, Device Error Code: 0x0000)
[No sense data]
This path is in DEAD state.
Path Recovery (3) has failed on device 0x3A path 0
(Adapter Error Code: 0x8007, Device Error Code: 0x0000)
[No sense data]
This path is in DEAD state.
END:datapath query device
Total Devices : 2
DEV#: 3A DEVICE NAME: 0x003A:[V596-A4-D1:0]
TYPE: 2105E20 SERIAL: 30812028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003A:[V596-A4-D1:0] DEAD NORMAL 1418 7
1 0x007A:[V596-A3-D1:0] OPEN NORMAL 4740 0
2 0x001A:[V596-A4-D0:0] DEAD NORMAL 1418 7
3 0x005A:[V596-A3-D0:0] OPEN NORMAL 4739 0
DEV#: 3B DEVICE NAME: 0x003B:[V596-A4-D1:1]
TYPE: 2105E20 SERIAL: 01312028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003B:[V596-A4-D1:1] OPEN NORMAL 1 0
1 0x007B:[V596-A3-D1:1] OPEN NORMAL 1 0
2 0x001B:[V596-A4-D0:1] OPEN NORMAL 1 0
3 0x005B:[V596-A3-D0:1] OPEN NORMAL 1 0
(If cable is continually disconnected, path will be set to INVALID if path
reclamation fails)
Path Reclamation has failed on device 0x3A path 2
(Adapter Error Code: 0x8007, Device Error Code: 0x0000)
[No sense data]
This path is in INVALID state.
Path Reclamation has failed on device 0x3A path 0
(Adapter Error Code: 0x8007, Device Error Code: 0x0000)
[No sense data]
This path is in INVALID state.
END:datapath query device
Total Devices : 2
DEV#: 3A DEVICE NAME: 0x003A:[V596-A4-D1:0]
TYPE: 2105E20 SERIAL: 30812028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003A:[V596-A4-D1:0] INVALID NORMAL 1418 8
1 0x007A:[V596-A3-D1:0] OPEN NORMAL 54740 0
2 0x001A:[V596-A4-D0:0] INVALID NORMAL 1418 8
3 0x005A:[V596-A3-D0:0] OPEN NORMAL 54739 0
DEV#: 3B DEVICE NAME: 0x003B:[V596-A4-D1:1]
TYPE: 2105E20 SERIAL: 01312028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003B:[V596-A4-D1:1] OPEN NORMAL 1 0
1 0x007B:[V596-A3-D1:1] OPEN NORMAL 1 0
2 0x001B:[V596-A4-D0:1] OPEN NORMAL 1 0
3 0x005B:[V596-A3-D0:1] OPEN NORMAL 1 0
(If pull both cable, volume will be deactivated, IO stops, paths will be set to
INVALID except one path left OPEN)
Aug 8, 2003 3:05:05 am NSS <comn>-3.02-xxxx: comnVol.c[7478]
Volume TEMPVOL: User data I/O error 20204(zio.c[1912]).
Block 268680(file block 63)(ZID 3779)

```

```

Volume TEMPVOL: User data I/O error 20204(zio.c[1912]).
Block 268681(file block 64)(ZID 3779)
Deactivating pool "TEMPPPOOL"...
Aug 8, 2003 3:05:06 am NSS<COMN>-3.02-xxxx: comnPool.c[2516]
Pool TEMPPPOOL: System data I/O error 20204(zio.c[1890]).
Block 610296(file block 10621)(ZID 3)
Dismounting Volume TEMPVOL
The share point "TEMPVOL" has been deactivated due to dismount of volume TEMPVOL
.
Aug 8, 2003 3:05:06 am NSS<COMN>-3.02-xxxx: comnVol.c[7478]
Volume TEMPVOL: User data I/O error 20204(zio.c[1912]).
Block 268682(file block 65)(ZID 3779)
Aug 8, 2003 3:05:07 am NSS<COMN>-3.02-xxxx: comnVol.c[7478]
Volume TEMPVOL: User data I/O error 20204(zio.c[1912]).
Block 268683(file block 66)(ZID 3779)
Aug 8, 2003 3:05:08 am NSS<COMN>-3.02-xxxx: comnVol.c[7478]
Block 268684(file block 67)(ZID 3779)
Aug 8, 2003 3:05:08 am NSS<COMN>-3.02-xxxx: comnVol.c[7478]
Block 268685(file block 68)(ZID 3779)
.....
END:datapath query device
Total Devices : 2
DEV#: 3A DEVICE NAME: 0x003A:[V596-A4-D1:0]
TYPE: 2105E20 SERIAL: 30812028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003A:[V596-A4-D1:0] OPEN NORMAL 2249 3064
1 0x007A:[V596-A3-D1:0] INVALID OFFLINE 12637 1
2 0x001A:[V596-A4-D0:0] INVALID OFFLINE 2248 16
3 0x005A:[V596-A3-D0:0] INVALID OFFLINE 12637 4
DEV#: 3B DEVICE NAME: 0x003B:[V596-A4-D1:1]
TYPE: 2105E20 SERIAL: 01312028 POLICY: Round Robin

Path# Device State Mode Select Errors
0 0x003B:[V596-A4-D1:1] OPEN NORMAL 1 0
1 0x007B:[V596-A3-D1:1] OPEN NORMAL 1 0
2 0x001B:[V596-A4-D0:1] OPEN NORMAL 1 0
3 0x005B:[V596-A3-D0:1] OPEN NORMAL 1 0
END:datapath query adapter
Active Adapters :2
Adpt# Adapter Name State Mode Select Errors Paths Active
0 [V596-A4] DEGRAD ACTIVE 4499 3080 4 2
1 [V596-A3] DEGRAD ACTIVE 25276 5 4 2
(After reconnect both cables, issue manual online command)
END:datapath set adapter 0 online
Success: set adapter 0 to online
Adpt# Adapter Name State Mode Select Errors Paths Active
0 [V596-A4] NORMAL ACTIVE 4499 3080 4 4
END:datapath set adapter 1 online
Success: set adapter 1 to online
Adpt# Adapter Name State Mode Select Errors Paths Active
1 [V596-A3] NORMAL ACTIVE 25276 5 4 4
END:datapath query adapter
Active Adapters :2
Adpt# Adapter Name State Mode Select Errors Paths Active
0 [V596-A4] NORMAL ACTIVE 4499 3080 4 4
1 [V596-A3] NORMAL ACTIVE 25276 5 4 4
(At this time, volume tempvol could not be mounted, pool activation is need)
END:mount tempvol
Volume TEMPVOL could NOT be mounted. Some or all volumes segments cannot be located.
If this is an NSS volume, the pool may need to be activated using the command nss /poolactivate=temppool.
END:nss /poolactivate=temppool
Activating pool "TEMPPPOOL"...
** Pool layout v40.07
** Processing journal
** 1 uncommitted transaction(s)
** 1839 Redo(s), 2 Undo(s), 2 Logical Undo(s)
** System verification completed
** Loading system objects
** Processing volume purge log
** .
** Processing pool purge log
** .
Loading volume "TEMPVOL"
Volume TEMPVOL set to the DEACTIVATE state.
Pool TEMPPPOOL set to the ACTIVATE state.
END:mount tempvol
Activating volume "TEMPVOL"
** Volume layout v35.00
** Volume creation layout v35.00
** Processing volume purge log
** .
Volume TEMPVOL set to the ACTIVATE state.
Mounting Volume TEMPVOL
** TEMPVOL mounted successfully
END:volumes
Mounted Volumes Name Spaces Flags

```

```
SYS DOS, LONG Cp Sa
 ADMIN DOS, MAC, NFS, LONG NSS P
TEMPVOL DOS, MAC, NFS, LONG NSS
3 volumes mounted
```

Chapter 7. Using the SDD on a Solaris host system

This chapter provides step-by-step procedures on how to install, configure, remove, and use SDD on a Solaris host system that is attached to supported storage devices. For updated and additional information not included in this manual, see the Readme file on the CD-ROM or visit the following SDD website:

www-01.ibm.com/support/docview.wss?uid=ssg1S4000052

Verifying the hardware and software requirements

You must install the following hardware and software components to ensure that SDD installs and operates successfully.

Hardware

The following hardware components are needed:

- One or more supported storage devices.
- For parallel SCSI access to ESS, one or more SCSI host adapters.
- One or more fibre-channel host adapters. In case of a single fibre-channel adapter, it must connect through a switch to multiple disk storage system ports.
- Subsystem LUNs that are created and confirmed for multiport access. Each LUN should have up to eight disk instances, with one for each path on the server.
- A SCSI cable to connect each SCSI host adapter to a storage system control-unit image port
- A fiber-optic cable to connect each fibre-channel adapter to a disk storage system controller port or a fibre-channel switch connected with disk storage system or virtualization product port. To install SDD and use the input/output (I/O) load-balancing and failover features, you need a minimum of two SCSI (ESS only) or fibre-channel host adapters if you are attaching to a disk storage system.

To install SDD and use the input-output (I/O) load-balancing and failover features, you need a minimum of two fibre-channel host adapters if you are attaching to a virtualization product.

SDD requires enabling the host-adapter persistent binding feature to have the same system device names for the same LUNs.

Software

SDD supports the following software components:

- ESS on a SPARC system running 32-bit Solaris 7/8/9 or 64-bit Solaris 7/8/9/10
- DS8000 on a SPARC system running 32-bit Solaris 8/9 or 64-bit Solaris 8/9/10
- DS8000 on an X64 machine running 64-bit Solaris 10
- DS6000 on a SPARC system running 32-bit Solaris 8/9 or 64-bit Solaris 8/9/10
- SAN Volume Controller on a SPARC system running 64-bit Solaris 8/9

SDD does not support the following software:

- Applications that issue SCSI 2 Reservation to storage

Supported environments

SDD supports 32-bit applications on 32-bit Solaris 7/8/9.

SDD supports both 32-bit and 64-bit applications on 64-bit Solaris 7/8/9/10.

Unsupported environments

SDD does not support the following environments:

- A host system with both a SCSI and fibre-channel connection to a shared LUN
- A system start from an SDD pseudo device
- A system paging file on an SDD pseudo device
- Root (/), /var, /usr, /opt, /tmp and swap partitions on an SDD pseudo device
- Single-path mode during concurrent download of licensed machine code nor during any disk storage system concurrent maintenance that impacts the path attachment, such as an disk storage system host-bay-adapter replacement
- Single-path configuration for Fibre Channel
- DS8000 and DS6000 do not support SCSI connectivity

Understanding how SDD works on a Solaris host system

SDD resides above the Solaris SCSI disk driver (sd) in the protocol stack. For more information about how SDD works, see “The SDD architecture” on page 2.

Preparing for SDD installation

Before you install SDD, you must first configure the disk storage systems or virtualization products to your host system.

Solaris SDD supports up to 600 physical LUNs from different storage systems on a single Solaris host.

Solaris SDD supports up to 32 paths per LUN.

Configuring disk storage systems

SDD requires a minimum of two independent paths that share the same logical unit to use the load-balancing and path-failover-protection features. **With a single path, failover protection is not provided.**

For information about how to configure your disk storage system, see the *IBM TotalStorage Enterprise Storage Server: Introduction and Planning Guide*.

Configuring virtualization products

Before you install SDD, configure your virtualization product and fibre-channel switches to assign LUNs to the system with multipath access. SDD requires a minimum of two independent paths that share the same logical unit to use the load-balancing and path-failover-protection features.

For information about configuring your SAN Volume Controller, see the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

Determining if the SDD server for Expert is installed

Note: IBM does not support ESS. For more information, contact the IBM Support team.

If you previously installed the SDD server (the stand-alone version) for IBM TotalStorage Expert V2R1 (ESS Expert) on your Solaris host system, you must remove this stand-alone version of the SDD server before you proceed with SDD 1.3.1.0 (or later) installation. The installation package for SDD 1.3.1.0 includes the SDD server daemon (also referred to as `sddsrv`), which incorporates the functionality of the stand-alone version of the SDD server (for ESS Expert).

To determine if the stand-alone version of the SDD server is installed on your host system, enter:

pkginfo -i SDDsrv

If you previously installed the stand-alone version of the SDD server, the output from the **pkginfo -i SDDsrv** command looks similar to the following output:

```
application SDDsrv SDDsrv bb-bit Version: 1.0.0.0 Nov-14-2001 15:34
```

Note:

- The installation package for the stand-alone version of the SDD server (for ESS Expert) is `SDDsrvSUNbb_yymmdd.pkg`. In this version, *bb* represents 32 or 64 bit, and *yymmdd* represents the date of the installation package. For ESS Expert V2R1, the stand-alone SDD server installation package is `SDDsrvSun32_020115.pkg` for a 32-bit environment and `SDDsrvSun64_020115.pkg` for a 64-bit environment.
- For instructions on how to remove the stand-alone version of the SDD server (for ESS Expert) from your Solaris host system, see the IBM® Subsystem Device Driver Server 1.0.0.0 (`sddsrv`) README for IBM TotalStorage Expert V2R1 at the following website:
www.ibm.com/servers/storage/support/software/swexpert/

For more information about the SDD server daemon, go to “SDD server daemon” on page 327.

Planning for installation

Before you install SDD on your Solaris host system, you need to understand what kind of software is running on it. The way that you install SDD depends on the kind of software that you are running. Three types of software communicate directly to raw or block disk-device interfaces such as `sd` and SDD:

- UNIX file systems, where no logical volume manager (LVM) is present.
- LVMs such as Sun Solstice Disk Suite. LVMs allow the system manager to logically integrate, for example, several different physical volumes to create the image of a single large volume.
- Major application packages, such as certain database managers (DBMSs).

You can install SDD in four different ways. The way that you choose depends on the kind of software that you have installed:

Table 22 further describes the various installation scenarios and how you should proceed.

Table 22. SDD installation scenarios

Installation scenario	Description	How to proceed
Scenario 1	<ul style="list-style-type: none"> • SDD is not installed. • No volume managers are installed. • No software application or DBMS is installed that communicates directly to the sd interface. 	Go to: <ol style="list-style-type: none"> 1. "Installing SDD" on page 315 2. "Standard UNIX applications" on page 329
Scenario 2	<ul style="list-style-type: none"> • SDD is not installed. • An existing volume manager, software application, or DBMS is installed that communicates directly to the sd interface. 	Go to: <ol style="list-style-type: none"> 1. "Installing SDD" on page 315 2. "Using applications with SDD" on page 329
Scenario 3	SDD is installed.	Go to "Upgrading the SDD" on page 320.
Scenario 4	Installing SDD CD-ROM in Jumpstart Environment	Go to "Installing SDD" on page 315.

Table 23 lists the installation package file names that come with SDD.

Table 23. Operating systems and SDD package file names

Operating system	Package file name
32-bit Solaris 7/8/9	sun32bit/IBMsdd
64-bit Solaris 7/8/9	sun64bit/IBMsdd
64-bit Solaris 10	solaris10/IBMsdd

For SDD to operate properly, ensure that the Solaris patches are installed on your operating system. Go to the following website for the latest information about Solaris patches:

www.oracle.com/technetwork/systems/patches/solaris/index.html

For more information on the Solaris patches, see the *IBM TotalStorage Enterprise Storage Server Host Systems Attachment Guide* or the *IBM System Storage SAN Volume Controller Host Systems Attachment Guide*.

Attention: Analyze and study your operating system and application environment to ensure that there are no conflicts with these patches prior to their installation.

Installing SDD

Before you install SDD, make sure that you have root access to your Solaris host system and that all the required hardware and software is ready. You can download the latest SDD package and readme from the SDD website:

www-01.ibm.com/support/docview.wss?uid=ssg1S4000052

Note: Note that SDD package name has changed from IBMdpo to IBMsdd for SDD 1.4.0.0 or later.

Installing SDD from CD-ROM

To install SDD on your Solaris host system from CD-ROM, complete the following steps.

Note: If the OS is Solaris 8 or Solaris 9, you can check the OS bit-level that is running by issuing `# isainfo -kv`.

1. Make sure that the SDD CD-ROM is available.
2. Insert the CD-ROM into your CD-ROM drive.
3. Change to the installation directory:

```
# cd /cdrom/cdrom0/sun32bit or  
# cd /cdrom/cdrom0/sun64bit
```

4. Issue the `pkgadd` command and point the `-d` option of the `pkgadd` command to the directory that contains IBMsdd. For Solaris 10, include the `-G` option to prevent SDD to be installed in nonglobal zones. For example,

```
pkgadd -d /cdrom/cdrom0/sun32bit IBMsdd or  
pkgadd -d /cdrom/cdrom0/sun64bit IBMsdd or  
pkgadd -G -d . /cdrom/cdrom0/Solaris10 IBMsdd
```

5. A message similar to the following message is displayed:

```
Processing package instance <IBMsdd> from <var/spool/pkg>
```

```
IBM SDD driver  
(sparc) 1  
## Processing package information.  
## Processing system information.  
## Verifying disk space requirements.  
## Checking for conflicts with packages already installed.  
## Checking for setuid/setgid programs.
```

```
This package contains scripts that run with super-user  
permission during the process of installing this package.
```

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

6. Enter `y` and press **Enter** to proceed. A message similar to the following message is displayed:

```

Installing IBM sdd driver as <IBMsdd>

## Installing part 1 of 1.
/etc/defvpath
/etc/rcS.d/S20vpath-config
/etc/sample_sddsrv.conf
/kernel/drv/sparcv9/vpathdd
/kernel/drv/vpathdd.conf
/opt/IBMsdd/bin/cfgvpath
/opt/IBMsdd/bin/datapath
/opt/IBMsdd/bin/defvpath
/opt/IBMsdd/bin/get_root_disk
/opt/IBMsdd/bin/pathstest
/opt/IBMsdd/bin/rmvpath
/opt/IBMsdd/bin/setlicense
/opt/IBMsdd/bin/showvpath
/opt/IBMsdd/bin/vpathmkdev
/opt/IBMsdd/devlink.vpath.tab
/opt/IBMsdd/etc.profile
/opt/IBMsdd/etc.system
/opt/IBMsdd/vpath.msg
/opt/IBMsdd/vpathexcl.cfg
/sbin/sddsrv
/usr/sbin/vpathmkdev
[ verifying class ]
## Executing postinstall script.

/etc/rcS.d/S20vpath-config
/etc/sample_sddsrv.conf
/kernel/drv/sparcv9/vpathdd
/kernel/drv/vpathdd.conf
/opt/IBMsdd/bin/cfgvpath
/opt/IBMsdd/bin/datapath
/opt/IBMsdd/bin/defvpath
/opt/IBMsdd/bin/get_root_disk
/opt/IBMsdd/bin/pathstest
/opt/IBMsdd/bin/rmvpath
/opt/IBMsdd/bin/setlicense
/opt/IBMsdd/bin/showvpath
/opt/IBMsdd/bin/vpathmkdev
/opt/IBMsdd/devlink.vpath.tab
/opt/IBMsdd/etc.profile
/opt/IBMsdd/etc.system
/opt/IBMsdd/vpath.msg
/opt/IBMsdd/vpathexcl.cfg
/sbin/sddsrv
/usr/sbin/vpathmkdev
[ verifying class ]
Vpath: Configuring 24 devices (3 disks * 8 slices)

Installation of <IBMsdd> was successful.

The following packages are available:
1 IBMcli ibm2105cli
   (sparc) 1.1.0.0
2 IBMsdd IBM SDD driver Version: May-10-2000 16:51
   (sparc) 1
Select package(s) you wish to process (or 'all' to process
all packages). (default: all) [?,??,q]:

```

7. If the SDD installation package determines that the system requires reboot, a message similar to the following message will be displayed:

```

*** IMPORTANT NOTICE ***
This machine must now be rebooted in order to ensure
sane operation. Issue
    shutdown -y -i6 -g0
and wait for the "Console Login:" prompt.

```

Installing SDD from downloaded code

There are several SDD packages for different Solaris versions on the SDD website. The SDD packages are in tar archives. After you download the appropriate SDD package, you can untar the file by issuing the following command:

```
tar xvf IBMsdd*.tar.
```

A directory IBMsdd is created. The installable package is located under the IBMsdd directory.

Install the package using steps similar to those found in “Installing SDD” on page 315.

Installing SDD in Jumpstart Environment

Installation of SDD in a Jumpstart environment is supported by SDD 1.6.2.0 and later. From this version, you can specify the base directory to which you want to install SDD:

```
pkgadd -R <basedir> IBMsdd
```

Using this command, SDD will be installed with the directory specified by *basedir* as a root directory. In this type of installation, vpath devices will not be configured during installation. You will need to reboot the system. Vpath devices will be automatically configured after reboot.

To install SDD in a Jumpstart environment, add installation of SDD using **-R** option in the Jumpstart finish script.

Postinstallation

If you install SDD from a CD-ROM, you can now manually unmount the CD. Issue the **umount /cdrom** command from the root directory. Go to the CD-ROM drive and press the Eject button.

After you install SDD, you might need to reboot the system to ensure proper operation. The SDD installation package determines if reboot is required. SDD displays a message to inform you to reboot only if reboot is required.

SDD vpath devices are found in the `/dev/rdisk` and `/dev/dsk` directories. The SDD vpath device is named according to the SDD instance number. A device with an instance number 1 would be: `/dev/rdisk/vpath1a` where *a* denotes a slice. Therefore, `/dev/rdisk/vpath1c` would be instance 1 and slice 2. Similarly, `/dev/rdisk/vpath2c` would be instance 2 and slice 2.

After SDD is installed, the device driver resides above the Sun SCSI disk driver (sd) in the protocol stack. In other words, SDD now communicates to the Solaris device layer. The SDD software installation procedure installs a number of SDD components and updates some system files. Table 24, Table 25 on page 318, and Table 26 on page 319 list those components and files.

Table 24. SDD components installed for Solaris host systems

File	Location	Description
vpathdd	/kernel/drv	Device driver
vpathdd.conf	/kernel/drv	SDD config file
Executables	/opt/IBMsdd/bin	Configuration and status tools

Table 24. SDD components installed for Solaris host systems (continued)

sddgetdata	/opt/IBMsdd/bin	The SDD data collection tool for problem analysis
S65vpath_config (except Solaris 10)	/etc/rcS.d	Boot initialization script Note: 1. This script must come before other LVM initialization scripts. 2. Prior to SDD 1.6.2.0, this file was named S20vpath_config.
ibmsdd-init.xml	/var/svc/manifest/system	SMF service manifest for boot time initialization (only on Solaris 10)
ibmsddinit	/lib/svc/method	Boot initialization script used by ibmsdd-init.xml manifest (only on Solaris 10)
sddsrv	/sbin/sddsrv	SDD server daemon
sample_sddsrv.conf	/etc/sample_sddsrv.conf	Sample SDD server config file

Table 25. System files updated for Solaris host systems

File	Location	Description
/etc/system	/etc	Forces the loading of SDD
/etc/devlink.tab	/etc	Tells the system how to name SDD devices in /dev

Table 26. SDD commands and their descriptions for Solaris host systems

Command	Description
cfgvpath [-c]	<p>Configures SDD vpath devices using the following process:</p> <ol style="list-style-type: none"> 1. Scan the host system to find all devices (LUNs) that are accessible by the Solaris host. 2. Determine which devices (LUNs) are the <i>same</i> devices that are accessible through different paths. 3. Create configuration file <code>/etc/vpath.cfg</code> to save the information about devices. <ul style="list-style-type: none"> • With <code>-c</code> option: cfgvpath exits without initializing the SDD driver. The SDD driver will be initialized after reboot. This option is used to reconfigure SDD after a hardware reconfiguration. Note: cfgvpath -c updates the configuration file but does not update the kernel. To update the kernel, you need to reboot. • Without <code>-c</code> option: cfgvpath initializes the SDD device driver <code>vpathdd</code> with the information stored in <code>/etc/vpath.cfg</code> and creates SDD vpath devices <code>/devices/pseudo/vpathdd*</code> Note: cfgvpath without <code>-c</code> option should not be used after hardware reconfiguration because the SDD driver is already initialized with previous configuration information. Reboot is required to properly initialize the SDD driver with the new hardware configuration information.
cfgvpath -r	<p>Reconfigures SDD vpath devices if SDD vpath devices exist. See “Option 2: Dynamic reconfiguration” on page 321. If no SDD vpath devices exist, use cfgvpath without <code>-r</code> option.</p>
showvpath	<p>Lists all SDD vpath devices and their underlying disks.</p>
vpathmkdev	<p>Creates files <code>vpathMsN</code> in the <code>/dev/dsk/</code> and <code>/dev/rdisk/</code> directories by creating links to the pseudo-vpath devices <code>/devices/pseudo/vpathdd*</code> that are created by the SDD driver.</p> <p>Files <code>vpathMsN</code> in the <code>/dev/dsk/</code> and <code>/dev/rdisk/</code> directories provide block and character access to an application the same way as the <code>cxydzsn</code> devices created by the system.</p> <p>vpathmkdev runs automatically during SDD package installation. However, issue manually to update files <code>vpathMsN</code> after hardware reconfiguration.</p>
datapath	<p>SDD driver console command tool.</p>

Table 26. SDD commands and their descriptions for Solaris host systems (continued)

Command	Description
<code>rmvpath [-b] [all vpathname]</code> <code>rmvpath -ab</code>	Removes SDD vpath devices from the configuration. See “Option 2: Dynamic reconfiguration” on page 321.

If you are not using a volume manager, software application, or DBMS that communicates directly to the sd interface, the installation procedure is nearly complete. If you have a volume manager, software application, or DBMS installed that communicates directly to the sd interface, such as Oracle, go to “Using applications with SDD” on page 329 and read the information specific to the application that you are using.

Verifying the SDD installation

To verify the SDD installation, complete the following steps:

1. Add `/opt/IBMsdd/bin` to the path.
 - a. C shell: `setenv PATH /opt/IBMsdd/bin:$PATH`
 - b. Bourne Shell: `PATH=/opt/IBMsdd/bin:$PATH, export PATH`
 - c. Korn Shell: `export PATH=/opt/IBMsdd/bin:$PATH`

To verify that you successfully installed SDD, enter **datapath query device**. If the command runs, SDD is installed.

Upgrading the SDD

To upgrade the SDD without requiring a reboot:

1. Stop I/O activity on all SDD devices.
2. In order to upgrade SDD without requiring reboot, all the SDD devices need to be closed before uninstalling SDD.
 - a. Unmount all the filesystems mounted on SDD devices.
 - b. If using Volume Manager, ensure that all the vpath devices are in offline mode in the Volume Manager.
 - c. To verify that all the SDD devices closed, issue the **datapath query adapter** command. The active path count for all the adapters should be zero.
3. Uninstall SDD using the procedure in “Uninstalling the SDD” on page 326.
4. Install SDD using the procedure in “Installing SDD” on page 315.

If, for some reason, SDD devices were not closed before step 3 and step 4, the system will need to be rebooted after SDD is installed. If reboot is required, the following message will be printed at the end of installation:

```
*** IMPORTANT NOTICE ***
This machine must now be rebooted in order to ensure
sane operation. Issue
    shutdown -y -i6 -g0
and wait for the "Console Login:" prompt.
```

Configuring the SDD

Before you start the SDD configuration process, make sure that you have successfully configured the disk storage system or virtualization product to which your host system is attached and that the disk storage system or virtualization product is operational.

Beginning with SDD 1.6.2.0, SDD automatically configures new vpath devices after each reboot. In the earlier versions, SDD only configures vpath devices during installation or reconfigures vpath devices initiated by the user running the SDD reconfiguration utility. From SDD 1.6.2.0 and later, SDD will reconfigure vpath devices after each reboot. It will automatically add or remove vpath devices or paths based on the device attachment during boot up.

Changing an SDD hardware configuration

When adding or removing multiport SCSI devices from your system, you must reconfigure SDD to recognize the new devices. Before reconfiguring SDD, the system needs to first recognize the hardware change.

Note: When an SDD vpath device, *vpathN*, is created for a LUN, SDD will also create a binding between that SDD vpath name, *vpathN*, to that LUN. The binding will not be removed even after the LUN has been removed from the host. The binding allows the same SDD vpath device name, *vpathN*, to be assigned to the same LUN when it is reconnected to the host. In order to reuse an SDD vpath name for a new LUN, the binding needed to be removed before reconfiguring SDD.

Option 1: Reconfigure the system and reconfigure the SDD

Complete the following steps to reconfigure the system and to reconfigure SDD. Step 1 and step 2 of this process reconfigure the system for the hardware change and the remaining steps reconfigure SDD.

1. Shut down the system. If you have a console attached to your host, enter **shutdown -i0 -g0 -y** and press **Enter**. If you do not have a console attached to your host, enter **shutdown -i6 -g0 -y** and press **Enter** to shut down and reboot the system.
2. If you have a console attached to your host (that is, you entered **shutdown -i0 -g0 -y** in step 1), perform a configuration restart by entering **boot -r** and pressing **Enter** at the OK prompt.
3. For SDD 1.6.2.0 or later, SDD automatically reconfigures after reboot. For SDD prior to 1.6.2.0, you must complete the following steps to reconfigure the SDD.
4. Run the SDD utility to reconfigure SDD. Enter **cfgvpath -c** and press **Enter**.
5. Shut down the system. Enter **shutdown -i6 -g0 -y** and press **Enter**.
6. After the restart, change to the `/opt/IBMsdd/bin` directory by entering:
cd /opt/IBMsdd/bin
7. Enter **devfsadm** and press **Enter** to reconfigure all the drives.
8. Enter **vpathmkdev** and press **Enter** to create all the SDD vpath devices.

Option 2: Dynamic reconfiguration

If the system can recognize the hardware change without reboot, dynamic reconfiguration provides a way to automatically detect path configuration changes

without requiring a reboot. After the system has recognized the new hardware change, the following commands will reconfigure the SDD.

Tip: Before issuing the following SDD dynamic reconfiguration commands, issue the **showvpath** and **datapath query device** commands. Save a copy of the output of both commands so that the change in the SDD configuration after the dynamic reconfiguration can be easily verified.

1. Issue the **cfgvpath -r** command.

Note: If there are no existing SDD vpath devices, the **cfgvpath -r** command does not dynamically reconfigure new SDD vpath devices. Issue **cfgvpath** to configure new SDD vpath devices. Then issue **devfsadm -i vpathdd** and **vpathmkdev**.

This operation finds the current hardware configuration and compares it to the SDD vpath device configuration in memory and then works out a list of differences. It then issues commands to put the SDD vpath device configuration in memory up-to-date with the current hardware configuration. The **cfgvpath -r** operation issues these commands to the vpath driver:

- a. Add one or more SDD vpath devices.

If you are adding new SDD vpath devices, issue **devfsadm -i vpathdd** and **vpathmkdev**.

- b. Remove an SDD vpath device; this will fail if the device is busy.

- c. Add one or more paths to the SDD vpath device.

If the SDD vpath device changes from single path to multiple paths, the path selection policy of the SDD vpath device will be changed to load-balancing policy.

- d. Remove a path for an SDD vpath device. This deletion of the path will fail if the device is busy, but the path will be set to DEAD and OFFLINE.

Removing paths of an SDD vpath device or removing an SDD vpath device can fail if the corresponding devices are busy. In the case of a path removal failure, the corresponding path would be marked OFFLINE. In the case of SDD vpath device removal failure, all the paths of the SDD vpath device would be marked OFFLINE. All OFFLINE paths would not be selected for I/Os. However, the SDD configuration file would be modified to reflect the paths or SDD vpath devices. When the system is rebooted, the new SDD configuration would be used to configure SDD vpath devices.

2. Issue the **rmvpath** command to remove one or more SDD vpath devices.

- a. To remove all SDD vpath devices that are not busy, issue the following command:

```
# rmpvpath -all
```

- b. To remove one SDD vpath device if the SDD vpath device is not busy, issue the following command:

```
# rmpvpath vpathname
```

For example, **rmvpath vpath10** will remove vpath10.

- c. To remove SDD vpath devices if the SDD vpath devices are not busy and also to remove the bindings between SDD vpath device names and LUNs so that the removed SDD vpath device names can be reused for new devices, issue the following command:

```
# rmpvpath -b -all
```

or

```
# rmpvpath -b vpathname
```

- d. To remove all bindings associated with currently unconfigured vpath names so that all unconfigured SDD vpath device names can be reused for new LUNs, issue the following command:

```
# rmvpath -ab
```

Note: Option **-ab** does not remove any existing SDD vpath devices.

Solaris 10 ZFS support

SDD 1.6.5.0 or later supports Solaris 10 ZFS. You can use vpath devices in place of native system devices in cXtYdZsN format when you create ZFS pools. For example, you can issue the following commands:

```
zpool create vpool1 vpath1a vpath1b vpath2c
zpool create vpool2 mirror vpath3c vpath4c
```

For information about known ZFS issues and the minimum Solaris 10 update that is required, see the Solaris SDD readme.

Solaris 10 Zone support

SDD cannot be installed in nonglobal zones; however, SDD can be installed in a global zone. Then, you can create device special files to port the SDD vpath devices to a nonglobal zone.

SDD installation in a server with nonglobal zones

CAUTION:

When you install SDD in a global zone that has nonglobal zones installed, you must prevent SDD from being installed in the nonglobal zone.

To do this, specify option **-G** when you issue the **pkgadd** command.

Note: You are no longer required to specify Option **-G** when you install Solaris SDD 1.6.5.0 or later. SDD installs only in global zones.

```
#pkgadd -G -d IBMsdd
```

Access SDD vpath devices in nonglobal zones

You can port vpath devices that are created in global zones to nonglobal zones, so they can be accessed in the nonglobal zone environment.

To port vpath devices into a nonglobal zone, complete the following steps:

1. Log in as root in the global zone.
2. Find the zone path of the nonglobal zone where you want to port the vpath device.

In the following example, the zonename is `/zones/ngb_zone2` for zone `ngb_zone2`.

```
# zonecfg -z ngb_zone2 info
zonename: ngb_zone2
zonepath: /zones/ngb_zone2
autoboot: true
pool:
limitpriv:
inherit-pkg-dir:
  dir: /lib
inherit-pkg-dir:
  dir: /platform
```

```
inherit-pkg-dir:
  dir: /sbin
inherit-pkg-dir:
  dir: /usr
```

3. Find the major and minor number of the vpath device that you want to port to the nonglobal zone by using the **ls -lL** command.

In the following example, the major number of vpath device *vpath2c* is 271 and the minor number is 18.

```
# ls -lL /dev/rdisk/vpath2c
crw----- 1 root  sys      271, 18 Jan 16 15:25 /dev/rdisk/vpath2c
```

4. Create block and raw device special files in the nonglobal zone */dev/dsk* and */dev/rdisk* directories using the **mknod** command. The device special files are based on the major and minor numbers of the vpath devices.

Issue the **mknod** command to create a block device special file:

```
# cd /zones/ngb_zone2/dev/dsk
# mnod vpath2c b 271 18
# ls -l total 0 brw-r--r-- 1 root  root  271, 18 Jan 16 15:59 vpath2c
```

Issue the **mknod** command to create a raw device special file:

```
#cd /zones/ngb_zone2/dev/rdisk
#mnod vpath2c c 271 18
#ls -l total 0 crw-r--r-- 1 root  root  271, 18 Jan 16 16:00 vpath2c
```

5. In the nonglobal zone environment, the vpath devices are now available in the */dev/dsk* and */dev/rdisk* directories.

Dynamically changing the SDD path-selection policy algorithm

SDD 1.4.0.0 (or later) supports multiple path-selection policies and allows users to change the path-selection policies dynamically. The following path-selection policies are supported:

failover only (fo)

All I/O operations for the device are sent to the same (preferred) path until the path fails because of I/O errors. Then an alternate path is chosen for subsequent I/O operations. This policy does not attempt to perform load balancing among paths.

load balancing (lb)

The path to use for an I/O operation is chosen by estimating the load on the adapter to which each path is attached. The load is a function of the number of I/O operations currently in process. If multiple paths have the same load, a path is chosen at random from those paths. Load-balancing mode also incorporates failover protection.

Note: The load-balancing policy is also known as the *optimized policy*.

round robin (rr)

The path to use for each I/O operation is chosen at random from paths that were not used for the last I/O operation. If a device has only two paths, SDD alternates between the two.

The path-selection policy is set at the SDD device level. The default path-selection policy for an SDD device is load balancing. You can change the policy for an SDD device. SDD version 1.4.0.0 (or later) supports dynamic changing of the path-selection policy for SDD devices.

Before changing the path-selection policy, determine the active policy for the device. Issue **datapath query device N**, where *N* is the device number of the SDD vpath device, to show the current active policy for that device. The output should look similar to the following example:

```

DEV#: 2 DEVICE NAME: vpath1c TYPE: 2105800 POLICY: OPTIMIZED
SERIAL: 03B23922
=====
Path# Adapter H/W Path Hard Disk State Mode Select Error
0 /pci@8,700000/fibre channel@3 sd@1,0:c,raw CLOSE NORMAL 0 0
1 /pci@8,700000/fibre channel@3 sd@2,0:c,raw CLOSE NORMAL 0 0
2 /pci@8,600000/fibre channel@1 sd@1,0:c,raw CLOSE NORMAL 0 0
3 /pci@8,600000/fibre channel@1 sd@2,0:c,raw CLOSE NORMAL 0 0

```

datapath set device policy command

Use the **datapath set device policy** command to change the SDD path-selection policy dynamically.

See “datapath set device policy” on page 389 for more information about the **datapath set device policy** command.

Excluding LUNs from being configured by SDD

Use the following procedure to exclude LUNs from being configured by SDD:

1. Determine the LUN identifier(s) of the LUN(s) to be excluded. See “Determining the LUN identifier of a LUN” for information about how to determine LUN identifiers.
2. Add the list of LUN identifiers to the `/etc/vpathexcl.cfg` file. Add one LUN Identifier per line.
3. Run **cfgvpath -r** to reconfigure SDD.

Vpath devices associated with the LUN identifiers in the `/etc/vpathexcl.cfg` file will be removed.

Subsequent SDD configuration will always exclude the LUN associated with the LUN identifier listed in `/etc/vpathexcl.cfg`.

If you want to re-create the vpath device for a LUN that has been excluded, remove the LUN identifier from `/etc/vpathexcl.cfg`.

Determining the LUN identifier of a LUN

After the SDD configuration, issue the **showvpath -l** command. The LUN identifier will be shown under the serial number.

For example:

```
showvpath -l
```

```

vpath47:   Serial Number : 13014712201
           Lun Identifier: 6005076303FFC05900000000000002201
c4t0d1s2  /devices/pci@0,0/pci1022,7450@b/pci1077,101@4/sd@0,1:c,raw
c4t3d1s2  /devices/pci@0,0/pci1022,7450@b/pci1077,101@4/sd@3,1:c,raw
c5t0d1s2  /devices/pci@0,0/pci1022,7450@b/pci1077,101@4,1/sd@0,1:c,raw
c5t1d1s2  /devices/pci@0,0/pci1022,7450@b/pci1077,101@4,1/sd@1,1:c,raw

```

Special consideration during SDD upgrade:

During SDD upgrade, /etc/vpathexcl.cfg is replaced and the LUN exclusion list is lost. In order to retain the exclusion list after SDD upgrade:

1. Copy the existing /etc/vpathexcl.cfg to a new file, for example, /etc/vpathexcl.cfg.sav, **before** installing the new SDD package.
2. After installing the new SDD package, replace /etc/vpathexec.cfg with the saved file, /etc/vpathexcl.cfg.sav.
3. Issue **cfgvpath -r** again to exclude the LUNs.

Uninstalling the SDD

The following procedure explains how to uninstall an SDD. You must uninstall the current level of SDD before you upgrade to a newer level. Because the SDD package name has changed from IBMdpo to IBMsdd for SDD 1.4.0.0 (or later), uninstall SDD requires you to uninstall either the IBMdpo or the IBMsdd package.

Complete the following steps to uninstall SDD:

1. Unmount all file systems on SDD devices.
2. If you are using SDD with a database, such as Oracle, edit the appropriate database configuration files (database partition) to remove all the SDD devices.
3. Enter **# pkgrm IBMdpo** or **# pkgrm IBMsdd** and press **Enter** depending on the previous SDD package installed.

Attention: A number of different installed packages are displayed. Make sure that you specify the correct package to uninstall.

A message similar to the following message is displayed:

```
The following package is currently installed:
IBMsdd IBMsdd Driver 64-bit Version: 1.6.0.5 Oct-21-2004 19:36
(sparc) 1.6.0.5

Do you want to remove this package? [y,n,?,q] y
```

4. Enter **y** and press **Enter**. A message similar to the following message is displayed:

```
## Removing installed package instance <IBMsdd>

This package contains scripts that run with super-user
permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] y
```

5. Enter **y** and press **Enter**. A message similar to the following message is displayed:

```
## Verifying package dependencies.
## Processing package information.
## Executing preremove script.
## Removing pathnames in class <none>
usr/sbin/vpathmkdev
/sbin/sddsrv
/opt/IBMsdd/vpathexcl.cfg
/opt/IBMsdd/vpath.msg
/opt/IBMsdd/etc.system
/opt/IBMsdd/etc.profile
/opt/IBMsdd/devlink.vpath.tab
/opt/IBMsdd/bin
/opt/IBMsdd
/kernel/drv/vpathdd.conf
/kernel/drv/sparcv9/vpathdd
/etc/sample_sddsrv.conf
/etc/rcS.d/S20vpath-config
/etc/defvpath
## Updating system information.

Removal of <IBMsdd> was successful.
```

Attention: If you are not performing an SDD upgrade, you should now reboot the system. If you are in the process of upgrading SDD, you do not need to reboot at this point. You can reboot the system after installing the new SDD package.

Understanding SDD support for single-path configuration for disk storage system

SDD does not support concurrent download of licensed machine code in single-path mode.

SDD does support single-path SCSI or fibre-channel connection from your SUN host system to a disk storage system. It is possible to create a volume group or an SDD vpath device with only a single path. However, because SDD cannot provide single-point-failure protection and load balancing with a single-path configuration, you should not use a single-path configuration.

SDD server daemon

The SDD server (also referred to as `sddsrv`) is an integrated component of SDD 1.3.1.0 (or later). This component consists of a UNIX application daemon that is installed in addition to the SDD device driver. See Chapter 9, “Using the SDDDSM or the SDDPCM daemon,” on page 365 for more information about `sddsrv`.

Verifying if the SDD server has started

After you have installed SDD, verify that the SDD server (`sddsrv`) has automatically started by entering `ps -ef | grep sddsrv`

If the SDD server (`sddsrv`) has automatically started, the output will display the process number on which `sddsrv` has started.

If the SDD server has *not* started, go to “Starting the SDD server manually” on page 328.

Starting the SDD server manually

If the SDD server does not start automatically after you complete the SDD installation or you want to start it manually after stopping `sddsrv`, use the following process to start `sddsrv`:

1. Edit `/etc/inittab` and verify the `sddsrv` entry. For example:

```
srv:234:respawn:/sbin/sddsrv > /dev/null 2>&1
```
2. Save the file `/etc/inittab`.
3. Issue `init q`.
4. Follow the directions in “Verifying if the SDD server has started” on page 327 to confirm that the SDD server started successfully.

Changing to a different port number for the SDD server

To change to a different port number for the SDD server, see the instructions in “Changing the `sddsrv` or `pcmsrv` TCP/IP port number” on page 369.

Changing the retry count value when probing SDD server inquiries

Beginning with SDD version 1.6.4.x on Solaris, you can dynamically change the value of the retry count for probing inquiries. Before SDD version 1.6.4.x on Solaris, the retry count was statically compiled in the `sddsrv` binary and was set to 2.

The retry count specifies how many consecutive probing inquiries that `sddsrv` makes before a path is determined to be nonfunctional and is set to DEAD state. A retry count of 2 indicates that `sddsrv` will attempt three consecutive probing inquiries.

You must either create the `probe_retry` variable in the existing `sddsrv.conf` file or generate a new `sddsrv.conf` file by copying the `sample_sddsrv.conf` file to the `sddsrv.conf` file in the `/etc` directory.

Make sure that this variable is commented with the default value of **`probe_retry=2`**. To change the default value, uncomment and set `probe_retry` to a valid value. The valid range for `probe_retry` is from 2 to 5. If you set a value that is not valid, `sddsrv` uses the default value of 2. If you set a value greater than 5, `sddsrv` uses 5. If you set a value less than 2, `sddsrv` uses 2.

Stopping the SDD server

To stop the SDD server, complete the following steps:

1. Edit `/etc/inittab` and comment out the SDD server entry:

```
#srv:234:respawn:/sbin/sddsrv > /dev/null 2>&1
```
2. Save the file.
3. Issue `init q`.
4. Check if `sddsrv` is running by running `ps -ef | grep sddsrv`. If `sddsrv` is still running, issue `kill -9 pid` of `sddsrv`.

Using applications with SDD

If your system already has a volume manager, software application, or DBMS installed that communicates directly with the Solaris disk device drivers, you need to insert the new SDD device layer between the program and the Solaris disk device layer. You also need to customize the volume manager, software application, or DBMS in order to have it communicate with the SDD devices instead of the Solaris devices.

In addition, many software applications and DBMS need to control certain device attributes such as ownership and permissions. Therefore, you must ensure that the new SDD devices accessed by these software applications or DBMS have the same attributes as the Solaris sd devices that they replace. You need to customize the software application or DBMS to accomplish this.

This section describes how to use the following applications with SDD:

- Standard UNIX applications
- NFS
- Veritas Volume Manager
- Oracle
- Solaris Volume Manager

Standard UNIX applications

If you have not already done so, install SDD using the procedure in “Installing SDD” on page 315. After you install SDD, the device driver resides above the Solaris SCSI disk driver (sd) in the protocol stack. In other words, SDD now communicates to the Solaris device layer.

Standard UNIX applications, such as **newfs**, **fsck**, **mkfs**, and **mount**, which normally take a disk device or raw disk device as a parameter, also accept the SDD device as a parameter. Similarly, you can replace entries in files such as **vfstab** and **dfstab** (in the format of **cntndnsn**) by entries for the corresponding SDD **vpathNs** devices. Make sure that the devices that you want to replace are replaced with the corresponding SDD device. Issue the **showvpath** command to list all SDD devices and their underlying disks.

Installing the SDD on a NFS file server

The procedures in this section show how to install SDD for use with an exported file system (NFS file server).

Setting up NFS for the first time

Complete the following steps if you are installing exported file systems on SDD devices for the first time:

1. If you have not already done so, install SDD using the procedure in the “Installing SDD” on page 315 section.
2. Determine which SDD (vpathN) volumes that you will use as file system devices.
3. Partition the selected volumes using the Solaris format utility.
4. Create file systems on the selected SDD devices using the appropriate utilities for the type of file system that you will use. If you are using the standard Solaris UFS file system, enter the following command:

```
# newfs /dev/rdisk/vpathNs
```

In this example, *N* is the SDD device instance of the selected volume. Create mount points for the new file systems.

5. Install the file systems into the `/etc/fstab` directory. Click **yes** in the **mount at boot** field.
6. Install the file system mount points into the directory `/etc/exports` for export.
7. Restart the system.

Installing SDD on a system that already has the Network File System file server

Complete the following steps to install SDD on a system if you have the Network File System file server already configured to export file systems that reside on a multiport subsystem and to use SDD partitions instead of `sd` partitions to access file systems:

1. List the mount points for all currently exported file systems by looking in the `/etc/exports` directory.
2. Match the mount points found in step 1 with `sdisk` device link names (files named `/dev/(r)dsk/cntndn`) by looking in the `/etc/fstab` directory.
3. Match the `sd` device link names found in step 2 with SDD device link names (files named `/dev/(r)dsk/vpathN`) by issuing the **showvpath** command.
4. Make a backup copy of the current `/etc/fstab` file.
5. Edit the `/etc/fstab` file, replacing each instance of an `sd` device link named `/dev/(r)dsk/cntndn` with the corresponding SDD device link.
6. Restart the system.
7. Verify that each exported file system:
 - Passes the start time **fsck pass**
 - Mounts properly
 - Is exported and available to NFS clients

If a problem exists with any exported file system after you complete step 7, restore the original `/etc/fstab` file and restart to restore Network File System service. Then review your steps and try again.

Veritas Volume Manager

For these procedures, ensure that you have a copy of the *Veritas Volume Manager System Administrator's Guide* and the *Veritas Volume Manager Command Line Interface for Solaris* to use for reference. These publications can be found at the following website:

www.symantec.com

The latest Veritas Volume Manager does not allow you to disable DMP. In order to use SDD as a multipathing driver under Veritas Volume Manager, you must install an appropriate ASL (device-specific library) for a specific IBM storage device type. There are two types of ASL for each device type:

- One type of ASL allows DMP to be the multipathing devices for a IBM device type
- Another type of ASL enables DMP to work in passthru mode and allows SDD to become the multipath driver for that device type

Each device type requires an appropriate ASL for that device type.

ASL to enable DMP passthru mode for ESS is called libvxvpath.so, which is normally included in the Veritas Volume Manger package.

ASL to enable DMP passthru mode for SAN Volume Controller is called libxsvc.so, and you must install it as a separate package. The ASL package is available for downloading from the website: www.symantec.com

SDD supports ESS devices for Veritas Volume Manager 3.5 MP2 or later and SAN Volume Controller devices for Veritas Volume Manager 3.5 MP2 Point Patch 3.1 or later with appropriate ASLs for SAN Volume Controller devices from Veritas.

To initially install SDD with Veritas Volume Manager:

Case 1: Installing Veritas Volume Manager for the first time.

1. Install SDD using the procedure in “Installing SDD” on page 315, if you have not already done so.
2. Ensure that you have rebooted the system after SDD is installed.
3. Install the Veritas Volume Manager package.
4. Follow the procedure in the Veritas Volume Manager manual to create the rootdg disk group and other required groups. In Veritas Volume Manager, the ESS vpath devices will have names such as VPATH_SHARK0_0, VPATH_SHARK0_1, and so on. SAN Volume Controller vpath devices will have names such as VPATH_SANVC0_0, VPATH_SANVC0_1, and so on.

Case 2: Installing SDD with Veritas already installed.

1. Install SDD using the procedure in “Installing SDD” on page 315.
2. Ensure that you have rebooted the system after SDD is installed.
In Veritas Volume Manager, the ESS vpath devices will have names such as VPATH_SHARK0_0, VPATH_SHARK0_1, and so on. SAN Volume Controller vpath devices will have names such as VPATH_SANVC0_0, VPATH_SANVC0_1, and so on.

Note: Multipathing of ESS and SAN Volume Controller devices managed by DMP before SDD installed will be managed by SDD after SDD is installed.

Oracle

You must have superuser privileges to complete the following procedures. You also need to have Oracle documentation on hand. These procedures were tested with Oracle 8.0.5 Enterprise server with the 8.0.5.1 patch set from Oracle.

Installing an Oracle database for the first time

You can set up your Oracle database in one of two ways. You can set it up to use a file system or raw partitions. The procedure for installing your database can differ depending on the choice that you make.

Using a file system:

1. If you have not already done so, install SDD using the procedure that is described in “Installing SDD” on page 315.
2. Create and mount file systems on one or more SDD partitions. (Oracle recommends three mount points on different physical devices.)

3. Follow the *Oracle Installation Guide* for instructions on how to install to a file system. (During the Oracle installation, you will be asked to name three mount points. Supply the mount points for the file systems that you created on the SDD partitions.)

Using raw partitions:

Attention: If using raw partitions, make sure all the databases are closed before going further. Make sure that the ownership and permissions of the SDD devices are the same as the ownership and permissions of the raw devices that they are replacing. Do not use disk cylinder 0 (sector 0), which is the disk label. Using it corrupts the disk. For example, slice 2 on Sun is the whole disk. If you use this device without repartitioning it to start at sector 1, the disk label is corrupted.

In the following procedure, you will replace the raw devices with the SDD devices.

1. If you have not already done so, install SDD using the procedure outlined in the “Installing SDD” on page 315 section.
2. Create the Oracle software owner user in the local server `/etc/passwd` file. You must also complete the following related activities:
 - a. Complete the rest of the Oracle preinstallation tasks described in the *Oracle8 Installation Guide*. Plan to install Oracle8 on a file system that resides on an SDD partition.
 - b. Set up the `ORACLE_BASE` and `ORACLE_HOME` environment variables of the Oracle user to be directories of this file system.
 - c. Create two more SDD-resident file systems on two other SDD volumes. Each of the resulting three mount points should have a subdirectory named `oradata`. The subdirectory is used as a control file and redo log location for the installer’s default database (a sample database) as described in the *Installation Guide*. Oracle recommends using raw partitions for redo logs. To use SDD raw partitions as redo logs, create symbolic links from the three redo log locations to SDD raw device links that point to the slice. These files are named `/dev/rdisk/vpathNs`, where *N* is the SDD instance number, and *s* is the partition ID.
3. Determine which SDD (`vpathN`) volumes you will use as Oracle8 database devices.
4. Partition the selected volumes using the Solaris format utility. If Oracle8 is to use SDD raw partitions as database devices, be sure to leave sector 0/disk cylinder 0 of the associated volume unused. This protects UNIX disk labels from corruption by Oracle8.
5. Ensure that the Oracle software owner has read and write privileges to the selected SDD raw partition device files under the `/devices/pseudo` directory.
6. Set up symbolic links in the `oradata` directory under the first of the three mount points. See step 2. Link the database files to SDD raw device links (files named `/dev/rdisk/vpathNs`) that point to partitions of the appropriate size.
7. Install the Oracle8 server following the instructions in the *Oracle Installation Guide*. Be sure to be logged in as the Oracle software owner when you run the `orainst /m` command. Select the **Install New Product - Create Database Objects** option. Select **Raw Devices** for the storage type. Specify the raw device links set up in step 2 for the redo logs. Specify the raw device links set up in step 3 for the database files of the default database.

8. To set up other Oracle8 databases, you must set up control files, redo logs, and database files following the guidelines in the *Oracle8 Administrator's Reference*. Make sure any raw devices and file systems that you set up reside on SDD volumes.
9. Launch the sqlplus utility.
10. Issue the **create database** SQL command, specifying the control, log, and system data files that you have set up.
11. Issue the **create tablespace** SQL command to set up each of the temp, rbs, tools, and users database files that you created.
12. Issue the **create rollback segment** SQL command to create the three redo log files that you set. For the syntax of these three **create** commands, see the *Oracle8 Server SQL Language Reference Manual*.

Installing an SDD on a system that already has Oracle in place

The installation procedure for a new SDD installation differs depending on whether you are using a file system or raw partitions for your Oracle database.

If using a file system:

Complete the following procedure if you are installing SDD for the first time on a system with an Oracle database that uses a file system:

1. Record the raw disk partitions being used (they are in the cntndnsn format) or the partitions where the Oracle file systems reside. You can get this information from the `/etc/vfstab` file if you know where the Oracle files are. Your database administrator can tell you where the Oracle files are, or you can check for directories with the name `oradata`.
2. Complete the basic installation steps in the "Installing SDD" on page 315 section.
3. Change to the directory where you installed the SDD utilities. Issue the **showvpath** command.
4. Check the directory list to find a `cntndn` directory that is the same as the one where the Oracle files are. For example, if the Oracle files are on `c1t8d0s4`, look for `c1t8d0s2`. If you find it, you will know that `/dev/dsk/vpath0c` is the same as `/dev/dsk/clt8d2s2`. (SDD partition identifiers end in an alphabetical character from a-g rather than s0, s1, s2, and so forth). A message similar to the following message is displayed:

```
vpath1c
c1t8d0s2 /devices/pci@1f,0/pci@1/scsi@2/sd@1,0:c,raw
c2t8d0s2 /devices/pci@1f,0/pci@1/scsi@2,1/sd@1,0:c,raw
```

5. Use the SDD partition identifiers instead of the original Solaris identifiers when mounting the file systems.

If you originally used the following Solaris identifiers:

```
mount /dev/dsk/c1t3d2s4 /oracle/mp1
```

you now use the following SDD partition identifiers:

```
mount /dev/dsk/vpath2e /oracle/mp1
```

For example, assume that `vpath2c` is the SDD identifier.

Follow the instructions in *Oracle Installation Guide* for setting ownership and permissions.

If using raw partitions: Complete the following procedure if you have Oracle8 already installed and want to reconfigure it to use SDD partitions instead of sd partitions (for example, partitions accessed through `/dev/rdisk/cntndn` files).

All Oracle8 control, log, and data files are accessed either directly from mounted file systems or through links from the `oradata` subdirectory of each Oracle mount point set up on the server. Therefore, the process of converting an Oracle installation from `sdisk` to SDD has two parts:

- Change the Oracle mount points' physical devices in `/etc/fstab` from `sdisk` device partition links to the SDD device partition links that access the same physical partitions.
- Re-create any links to raw `sdisk` device links to point to raw SDD device links that access the same physical partitions.

Converting an Oracle installation from sd to SDD partitions:

Complete the following steps to convert an Oracle installation from `sd` to SDD partitions:

1. Back up your Oracle8 database files, control files, and redo logs.
2. Obtain the `sd` device names for the Oracle8 mounted file systems by looking up the Oracle8 mount points in `/etc/vfstab` and extracting the corresponding `sd` device link name (for example, `/dev/rdisk/c1t4d0s4`).
3. Launch the `sqlplus` utility.
4. Enter the command:

```
select * from sys.dba_data_files;
```

The output lists the locations of all data files in use by Oracle. Determine the underlying device where each data file resides. You can do this by either looking up mounted file systems in the `/etc/vfstab` file or by extracting raw device link names directly from the `select` command output.

5. Enter the `ls -l` command on each device link found in step 4 and extract the link source device file name. For example, if you enter the command:

```
# ls -l /dev/rdisk/c1t1d0s4
```

A message similar to the following message is displayed:

```
/dev/rdisk/c1t1d0s4 /devices/pci@1f,0/pci@1/scsi@2/sd@1,0:e
```

6. Write down the file ownership and permissions by issuing the `ls -lL` command on either the files in `/dev/` or `/devices` (it yields the same result). For example, if you enter the command:

```
# ls -lL /dev/rdisk/c1t1d0s4
```

A message similar to the following message is displayed:

```
crw-r--r-- oracle dba 32,252 Nov 16 11:49 /dev/rdisk/c1t1d0s4
```

7. Complete the basic installation steps in the “Installing SDD” on page 315 section.

8. Match each cntndns device with its associated vpathNs device link name by issuing the **showvpath** command. Remember that vpathNs partition names use the letters a - h in the s position to indicate slices 0 - 7 in the corresponding cntndns slice names.
9. Issue the **ls -l** command on each SDD device link.
10. Write down the SDD device nodes for each SDD device link by tracing back to the link source file.
11. Change the attributes of each SDD device to match the attributes of the corresponding disk device using the **chgrp** and **chmod** commands.
12. Make a copy of the existing /etc/vfstab file for recovery purposes. Edit the /etc/vfstab file, changing each Oracle device link to its corresponding SDD device link.
13. For each link found in an oradata directory, re-create the link using the appropriate SDD device link as the source file instead of the associated sd device link. As you perform this step, generate a reversing shell script that can restore all the original links in case of error.
14. Restart the server.
15. Verify that all file system and database consistency checks complete successfully.

Solaris Volume Manager (formerly Solstice DiskSuite)

Note: Sun has renamed Solstice DiskSuite to Solaris Volume Manager.

The following procedure applies to Solaris Volume Manager. Depending on the DiskSuite version, the md.tab file is in either the /etc/opt/SUNWmd/ directory or the /etc/lvm/ directory.

For these procedures, you need access to the Solaris answerbook facility. These procedures were tested using Solstice DiskSuite 4.2 with the patch 106627-04 (DiskSuite patch) installed. You should have a copy of the *DiskSuite Administration Guide* available to complete these procedures. You must have superuser privileges to complete these procedures.

Note: SDD only supports Solstice DiskSuite line command interface. The DiskSuite Tool (metatool) does not recognize and present SDD devices for configuration. SDD does not support Solaris Volume Manager disk set feature, which issues SCSI 2 Reservation to storage.

Installing Solaris Volume Manager for the first time

Complete the following steps if you are installing Solaris Volume Manager on the multiport subsystem server for the first time:

1. Install SDD using the procedure in the “Installing SDD” on page 315 section, if you have not already done so.
2. Use the **boot -r** command to configure the SPARC server to recognize all devices over all paths.
3. Install the Solaris Volume Manager packages and the answerbook. Do not restart yet.
4. Determine which SDD vpath devices that you will use to create Solaris Volume Manager metadevices. Partition these devices by selecting them in the Solaris format utility. The devices appear as vpathNs, where N is the vpath driver instance number. Use the partition submenu, just as you would for an sd device link of the form, cntndn. If you want to know which cntndn links

correspond to a particular SDD vpath device, enter the **showvpath** command and press **Enter**. Reserve at least three partitions of three cylinders each for use as Solaris Volume Manager Replica database locations.

Note: You do not need to partition any sd (cntndn) devices.

5. Set up the replica databases on a separate partition. This partition needs to be at least three partitions of three cylinders. Do not use a partition that includes Sector 0 for this database replica partition. Complete the following instructions for setting up replica databases on the vpathNs partitions, where *N* is the SDD vpath device instance number and *s* is the letter denoting the three-cylinder partition, or slice, of the device that you want to use as a replica. Remember that partitions a - h of an SDD vpath device correspond to slices 0 - 7 of the underlying multiport subsystem device.

Note: Verify that Solaris Volume Manager on the host supports replica databases on SAN devices before you set up replica databases on SDD vpath devices.

6. Follow the instructions in the *Solaris Volume Manager Administration Guide* to build the types of metadevices that you need. Use the **metainit** command and the /dev/(r)dsk/vpathNs device link names wherever the instructions specify /dev/(r)dsk/cntndnsn device link names.
7. Insert the setup of all vpathNs devices that are used by DiskSuite into the md.tab file.

Installing SDD on a system that already has Solstice DiskSuite in place

Complete the following steps if Solstice DiskSuite is already installed and you want to convert existing sd devices used in metadevice configuration to the corresponding SDD devices:

1. Back up all data.
2. Back up the current Solstice configuration by making a copy of the md.tab file and recording the output of the **metastat** and **metadb -i** commands. Make sure all sd device links in use by DiskSuite are entered in the md.tab file and that they all come up properly after a restart.
3. Install SDD using the procedure in the “Installing SDD” on page 315 section, if you have not already done so. After the installation completes, enter **shutdown -i6 -y -g0** and press **Enter**. This verifies the SDD vpath installation.

Note: Do not do a reconfiguration restart.

4. Using a plain sheet of paper, make a two-column list and match the /dev/(r)dsk/cntndnsn device links found in step 2 with the corresponding /dev/(r)dsk/vpathNs device links. Use the **showvpath** command to do this step.
5. Delete each replica database that is currently configured with a /dev/(r)dsk/cntndnsn device by using the **metadb -d -f <device>** command. Replace the replica database with the corresponding /dev/(r)dsk/vpathNs device found in step 2 by using the **metadb -a <device>** command.
6. Create a new md.tab file. Insert the corresponding vpathNs device link name in place of each cntndnsn device link name. Do not do this for start device partitions (vpath does not currently support these). When you are confident that the new file is correct, install it in either the /etc/opt/SUNWmd directory or the /etc/lvm directory, depending on the DiskSuite version.
7. Restart the server, or proceed to the next step if you want to avoid restarting your system.

To back out the SDD vpath in case there are any problems following step 7:

- a. Reverse the procedures in step 4 to step 6, reinstalling the original md.tab in the /etc/opt/SUNWmd directory or the /etc/lvm directory depending on the DiskSuite version.
 - b. Enter the **pkgrm IBMsdd** command.
 - c. Restart.
8. Stop all applications using DiskSuite, including file systems.
 9. Enter the following commands for each existing metadvice:

metaclear <device>

10. Enter **metainit -a** to create metadevices on the /dev/(r)dsk/vpathNs devices.
11. Compare the metadevices that are created with the saved metastat output from step 2.
Create any missing metadevices and reconfigure the metadevices based on the configuration information from the saved metastat output.
12. Restart your applications.

Setting up transactional volume for UFS logging on a new system

For these procedures, you need access to the Solaris answerbook facility. You must have super-user privileges to complete these procedures.

Complete the following steps if you are installing a new UFS logging file system on SDD vpath devices:

1. Install SDD using the procedure in the “Installing SDD” on page 315 section, if you have not already done so.
2. Determine which SDD vpath (vpathNs) volumes that you will use as file system devices. Partition the selected SDD vpath volumes using the Solaris format utility. Be sure to create partitions for UFS logging devices as well as for UFS master devices.
3. Create file systems on the selected vpath UFS master device partitions using the **newfs** command.
4. Install Solaris Volume Manager if you have not already done so.
5. Create the metatrans device using **metainit**. For example, assume /dev/dsk/vpath1d is your UFS master device used in step 3, /dev/dsk/vpath1e is its corresponding log device, and d0 is the trans device that you want to create for UFS logging. Enter **metainit d0 -t vpath1d vpath1e** and press **Enter**.
6. Create mount points for each UFS logging file system that you have created using steps 3 and 5.
7. Install the file systems into the /etc/vfstab directory, specifying /dev/md/(r)dsk/d <metadvice number> for the raw and block devices. Set the **mount at boot** field to *yes*.
8. Restart your system.

Installing vpath on a system that already has transactional volume for UFS logging in place

Complete the following steps if you already have UFS logging file systems residing on a multiport subsystem and you want to use vpath partitions instead of sd partitions to access them.

1. Make a list of the DiskSuite metatrans devices for all existing UFS logging file systems by looking in the /etc/vfstab directory. Make sure that all configured

metatrans devices are set up correctly in the md.tab file. If the devices are not set up now, set them up before continuing. Save a copy of the md.tab file.

2. Match the device names found in step 1 with sd device link names (files named /dev/(r)dsk/cntndnsn) using the **metastat** command.
3. Install SDD using the procedure in the "Installing SDD" on page 315 section, if you have not already done so.
4. Match the sd device link names found in step 2 with SDD vpath device link names (files named /dev/(r)dsk/vpathNs) by issuing the **/opt/IBMsdd/bin/showvpath** command.
5. Unmount all current UFS logging file systems known to reside on the multiport subsystem using the **umount** command.
6. Enter **metaclear -a** and press **Enter**.
7. Create new metatrans devices from the vpathNs partitions found in step 4 that correspond to the sd device links found in step 2. Remember that vpath partitions a - h correspond to sd slices 0 - 7. Use the **metainit d <metadevice number> -t <"vpathNs" - master device> <"vpathNs" - logging device>** command. Be sure to use the same metadevice numbering as you originally used with the sd partitions. Edit the md.tab file to change each metatrans device entry to use vpathNs devices.
8. Restart the system.

Note: If there is a problem with a metatrans device after steps 7 and 8, restore the original md.tab file and restart the system. Review your steps and try again.

Chapter 8. Using SDDDSM on a Windows Server 2008 or Windows Server 2012 host system

Subsystem Device Driver Device Specific Module (SDDDSM) provides multipath I/O support based on the MPIO technology of Microsoft. SDDDSM is a device-specific module designed to provide support for supported storage devices.

This chapter provides procedures for you to install, configure, use, and remove SDDDSM on a Windows Server 2008, or Windows Server 2012 host system that is attached to a supported storage device.

Install the package from the %ProgramFiles%\IBM\SDDDSM directory of the SDDDSM CD-ROM or the location where SDDDSM package was saved.

For updated and additional information that is not included in this chapter, see the readme file on the CD-ROM or visit the SDDDSM website:

www-01.ibm.com/support/docview.wss?uid=ssg1S4000350

Verifying the hardware and software requirements

You must have the following hardware and software components in order to install SDDDSM:

Hardware

The following hardware components are needed:

- Supported storage devices
- Host system
- Fibre-channel adapters and cables

Software

The following software components are needed:

- Operating system:
 - Windows Server 2008 (32-bit)
 - Windows Server 2008, or Windows Server 2012 (x64)
- Device driver for fibre-channel adapters (it must be a STORPORT-based miniport driver)
- SDDDSM installation package, which includes:
 - MPIO drivers. MPIO is shipped with Windows Server 2008 and later. On Windows Server 2008 and Windows Server 2012, the MPIO drivers that are shipped with the operating system are used.
 - SDDDSM driver: sdddsm.sys
 - SDDDSM installer: setup.exe
 - SDDDSM readme file
 - SDDDSM command windows shortcut
- The Microsoft Visual C++ 2012 Redistributable package that can be downloaded from the Microsoft Corporation website. The HBAInfo utility requires this package.

Unsupported environments

SDDDSM does not support the following environments:

- Single-path mode during code distribution and activation of LMC nor during any disk storage system concurrent maintenance that impacts the path attachment, such as a disk storage system host-bay-adapter replacement.
- DS8000 and DS6000 do not support SCSI connectivity.
- SDDDSM is not supported with Windows Server 2003 from version 2.4.5.0 onwards.

Host system requirements

To install SDDDSM, your Windows Server 2008 host must be an x86-based 32-bit system or an AMD-based 64-bit system. The Windows Server 2012 host system must be an AMD-based 64-bit system. The host system can be a uniprocessor system or a multiprocessor system.

To support SDDDSM, your Windows Server 2008 or Windows Server 2012 host system must have the following installed:

- Any Windows Server 2008 or Windows Server 2012 hotfixes, when applicable
- Any HBA-required Windows Server 2008 or Windows Server 2012 hotfixes, when applicable

For additional information, see the SDDDSM readme and the HBA driver readme or release notes.

To install all components, you must have at least 10 MB (MB equals approximately 1 000 000 bytes) of disk space available on the drive where Windows Server 2008 or Windows Server 2012 is installed.

Fibre-channel requirements

To use the SDDDSM fibre-channel support, ensure that your host system meets the following requirements:

- No more than 32 fibre-channel adapters are attached.
- A fiber-optic cable connects each fibre-channel adapter to a disk storage system port.
- If you need the SDDDSM I/O load-balancing and failover features, ensure that a minimum of two fibre-channel adapters are installed.

Note: If your host has only one fibre-channel adapter, it requires you to connect through a switch to multiple disk storage system ports. You should have at least two fibre-channel adapters to prevent data loss due to adapter hardware failure or software failure.

For information about the fibre-channel adapters that can attach to your Windows Server 2008 or Windows Server 2012 host system, go to the following website:

www.ibm.com/servers/storage/support

Preparing for SDDDSM installation

Before installing SDDDSM, you must:

1. Ensure that all hardware and software requirements are met. For more information, see “Verifying the hardware and software requirements” on page 339.

2. Configure the fibre-channel adapters that are attached to your Windows Server 2008 or Windows Server 2012 host system. See “Configuring fibre-channel adapters” for more information.
3. Review the SDDDSM uninstallation and installation instructions. For more information, see “Uninstalling SDDDSM” on page 350 and “Installing SDDDSM.”

Configuring the supported storage device

Before you install SDDDSM, configure your supported storage device for single-port or multiport access for each LUN. To use the load-balancing and failover features, SDDDSM requires a minimum of two independent paths that share the same logical unit.

For information about configuring your supported storage device, see the Implementation Guide or Configuration Guide for your product.

Configuring fibre-channel adapters

You must configure the fibre-channel adapters that are attached to your Windows host system before you install SDDDSM. Follow the adapter-specific configuration instructions to configure the adapters attached to your Windows Server 2008 or Windows Server 2012 host systems.

Installing SDDDSM

You can install SDDDSM either from a CD-ROM or download.

Installing for the first time

The following section describes how to install SDDDSM on your system.

Note: Ensure that SDDDSM is installed *before* adding new paths to a device. Otherwise, the Windows Server 2008 or Windows Server 2012 server could cause data corruption on that device.

Installing SDDDSM from CD-ROM

Complete the following steps to install SDDDSM on your system:

1. If this is a new installation, ensure that there is a single connection from your host to your storage. Multipath access to the storage can be added after SDDDSM is installed.
2. Log on as the administrator user.
If you plan to install SDDDSM from CD-ROM, continue with step 3; otherwise, go to step 9.
3. Insert the SDDDSM installation CD-ROM into the selected drive.
4. Start the Windows Server 2008 or Windows Server 2012 Explorer program.
5. Select the CD-ROM drive. A list of all the installed directories on the compact disc is displayed.
6. If you have the zip file for the SDDDSM package available, select the %ProgramFiles%\IBM\SDDDSM installation subdirectory and go to step 9.
7. If you still do not have the zip file for the SDDDSM package available, go to the SDD website and download and save it to a directory.
8. Extract the zip file for the SDDDSM package to a directory and go to that directory.

9. Select the 32-bit or 64-bit folder depending on the operating system of the host system.
10. Run the setup.exe program. Follow the instructions.
11. Shut down the host system.
12. Connect additional cables to your storage if needed.
13. Make any necessary zoning configuration changes.
14. Restart the host system.

After completing the installation procedures and when you log on again, you will see an **SDDDSM** entry in your Program menu containing the following selections:

1. Subsystem Device Driver DSM
2. SDDDSM Technical Support website
3. README

Note:

1. You can verify that SDDDSM has been successfully installed by issuing the **datapath query device** command. You must issue the **datapath** command from the datapath directory. If the command runs, SDDDSM is installed.

You can also use the following operation to verify that SDDDSM has been successfully installed:

- a. Click **Start → Programs → Administrative Tools → Computer Management**.
- b. Double-click **Device Manager**.
- c. Expand **Disk drives** in the right pane.

Note: In Windows 2012 Server Core CLI mode, run the **Get-Disk** command in Windows PowerShell to see the disk list.

In Figure 7, there are eight SAN Volume Controller devices connected to the host and four paths to each of the SAN Volume Controller devices. The Device manager shows eight 2145 Multipath Disk Devices and 32 2145 SDDDSM SCSI Devices.

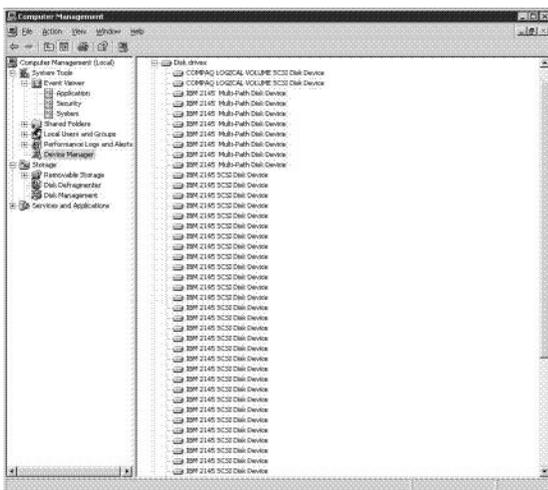


Figure 7. Example showing SAN Volume Controller devices to the host and path access to the SAN Volume Controller devices in a successful SDDDSM installation on a Windows Server host system

2. You can also verify the current version of SDDDSM. For more information, go to “Displaying the current version of SDDDSM” on page 343.

Installing SDDDSM from downloaded code

Use this procedure to install SDDDSM from downloaded code:

1. Unzip the SDDDSM code to your installation subdirectory.
2. Select the 32-bit or 64-bit folder depending on the operating system of the system.
3. Run the setup.exe program and follow the instructions. **Tip:** The setup.exe program offers the following command line options for silent install:
--> setup -s : silent install
--> setup -s -n : silent install; no reboot (requires SDDDSM 2.1.1.0 or later)
4. When the setup.exe program is finished, you will be asked if you want to reboot. If you answer **y**, the setup.exe program will restart your SDDDSM system immediately. Follow the instructions to restart. Otherwise, the setup.exe program exits, and you need to manually restart your SDDDSM system to activate the new installation.
5. Shut down your SDDDSM host system.
6. Reconnect all cables that connect the host bus adapters and the supported storage devices if needed.
7. Change any zoning information that needs to be updated.
8. Restart your SDDDSM host system.

Upgrading SDDDSM

Use the following procedure to upgrade SDDDSM.

1. Log in as administrator.
2. Open Windows Explorer and go to the directory where the SDDDSM package is located.
3. Run the setup.exe program and follow the instructions to continue with the SDDDSM setup procedure.
4. When the upgrade is complete, SDDDSM will ask you to reboot. Answer yes to reboot the system and activate the new SDDDSM.

You can check the SDDDSM version to verify that SDDDSM has been successfully upgraded. For more information, see “Displaying the current version of SDDDSM.”

Displaying the current version of SDDDSM

You can display the current version of SDDDSM on a Windows Server 2008 or Windows Server 2012 host system by viewing the sdddsm.sys file properties. Complete the following steps to view the properties of sdddsm.sys file:

1. Click **Start** → **Programs** → **Accessories** → **Windows Explorer** to open Windows Explorer.
2. In Windows Explorer, go to the %SystemRoot%\system32\drivers directory, where %SystemRoot% is:
%SystemDrive%\Windows for Windows Server 2008 or Windows Server 2012
If Windows is installed on the C: drive, %SystemDrive% is C:. If Windows is installed on the E: drive, %SystemDrive% is E:
3. Right-click the sdddsm.sys file, and then click **Properties**. The sdddsm.sys properties window opens.
4. In the sdddsm.sys properties window, click **Version**. The file version and copyright information about the sdddsm.sys file is displayed.

If the SDDDSM version is 2.1.1.0 or later, you can use the **datapath query version** command:

1. Open the SDDDSM command window.
2. Issue the **datapath query version** command.

Configuring SDDDSM

Use these topics to configure SDDDSM.

Maximum number of LUNs

There is no physical limitation on the maximum number of devices that can be configured for SDDDSM. Any limitation exists on the host operating system.

Consult the *Host Attachment Guide* for your product.

Verifying the configuration

To activate SDDDSM, you need to restart your Windows Server 2008 or Windows Server 2012 system after it is installed.

Attention: Ensure that SDDDSM is installed and activated before you add new paths to a device. Otherwise, the Windows Server 2008 or Windows Server 2012 server could cause data corruption on that device.

Before adding any new hardware, review the configuration information for the adapters and devices currently on your Windows Server 2008 or Windows Server 2012 server. Complete the following steps to display information about the adapters and devices:

1. You must log on as an administrator user.
2. Open the DOS prompt window.
 - On a Windows Server 2008: Click **Start** → **Program** → **Subsystem Device Driver DSM** → **Subsystem Device Driver DSM**.
 - On a Windows Server 2012: Click **Start Screen** → **Apps** → **Subsystem Device Driver DSM**.

On a Windows Server 2012 Server Core CLI: Go to %PROGRAMFILES%\IBM\SDDDSM.

3. Enter **datapath query adapter** and press **Enter**. The output includes information about all the installed adapters. In the following example, the output shows that one HBA is installed:

```
Active Adapters : 1
Adpt# Name           Interface State  Mode  Select Errors Paths  Active
0 Scsi Port6 Bus0 FC     NORMAL ACTIVE 1099  1    112   112
```

4. Enter **datapath query device** and press **Enter**. In the following example, the output shows that three devices are attached to the SCSI path:

Total Devices : 3

DEV#: 0 DEVICE NAME: Disk1 Part0 TYPE: 2145 POLICY: ROUND ROBIN SUB SET
SERIAL: 60050768019B02A810000000000001A1 Reserved: No LUN SIZE: 20.0GB
HOST INTERFACE: FC

Path#	Adapter/Hard Disk	State	Mode	Select	Errors
0 *	Scsi Port6 Bus0/Disk1 Part0	OPEN	NORMAL	51	1
1 *	Scsi Port6 Bus0/Disk1 Part0	OPEN	NORMAL	0	0
2	Scsi Port6 Bus0/Disk1 Part0	OPEN	NORMAL	0	0
3	Scsi Port6 Bus0/Disk1 Part0	OPEN	NORMAL	0	0

DEV#: 1 DEVICE NAME: Disk2 Part0 TYPE: 2145 POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A810000000000001A2 Reserved: No LUN SIZE: 20.0GB
HOST INTERFACE: FC

Path#	Adapter/Hard Disk	State	Mode	Select	Errors
0	Scsi Port6 Bus0/Disk2 Part0	OPEN	NORMAL	16	0
1	Scsi Port6 Bus0/Disk2 Part0	OPEN	NORMAL	9	0
2 *	Scsi Port6 Bus0/Disk2 Part0	OPEN	NORMAL	0	0
3 *	Scsi Port6 Bus0/Disk2 Part0	OPEN	NORMAL	0	0

DEV#: 2 DEVICE NAME: Disk3 Part0 TYPE: 2145 POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A810000000000001A3 Reserved: No LUN SIZE: 20.0GB
HOST INTERFACE: FC

Path#	Adapter/Hard Disk	State	Mode	Select	Errors
0 *	Scsi Port6 Bus0/Disk3 Part0	OPEN	NORMAL	14	0
1 *	Scsi Port6 Bus0/Disk3 Part0	OPEN	NORMAL	10	0
2	Scsi Port6 Bus0/Disk3 Part0	OPEN	NORMAL	0	0
3	Scsi Port6 Bus0/Disk3 Part0	OPEN	NORMAL	0	0

Activating additional paths

Complete the following steps to activate additional paths to an SDDDSM device:

1. Install any new hardware on the Windows Server 2008 or Windows Server 2012 server, or the SAN Volume Controller.
2. Open the DOS prompt window.
 - On a Windows Server 2008: Click **Start** → **Program** → **Administrative Tools** → **Computer Management**.
 - On a Windows Server 2012: Click **Server Manager** → **Tools** → **Computer Management**.
3. Click **Device Manager**.
4. Right-click **Disk drives**.
5. Click **Scan for hardware changes**.

Note: To obtain the latest information about disks in the Windows 2012 Server Core CLI, run the **Get-Disk** command in the Windows PowerShell.

6. Verify that the path is added correctly. See “Verifying that additional paths are installed correctly.”

Verifying that additional paths are installed correctly

After installing additional paths to SDDDSM devices, verify that the additional paths have been installed correctly.

Complete the following steps to verify that the additional paths are installed correctly:

1. Open the DOS prompt window.

- On a Windows Server 2008: Click **Start** → **Program** → **Subsystem Device Driver DSM** → **Subsystem Device Driver DSM**.
- On a Windows Server 2012: Click **Start** → **Apps** → **Subsystem Device Driver DSM**.

2. Type **datapath query adapter** and press **Enter**. The output includes information about any additional adapters that were installed. In the example shown in the following output, an additional HBA has been installed:

```
Active Adapters :2
Adpt#      Name      Interface  State  Mode  Select  Errors  Paths  Active
   0  Scsi Port6 Bus0      FC    NORMAL  ACTIVE  1099    1     112    112
   1  Scsi Port7 Bus0      FC    NORMAL  ACTIVE   813    0     112    112
```

3. Type **datapath query device** and press **Enter**. The output should include information about any additional devices that were installed. In this example, the output includes information about the new HBA and the new device numbers that were assigned. The following output is displayed:

```
Total Devices : 3

DEV#:  0  DEVICE NAME: Disk1 Part0      TYPE: 2145  POLICY: ROUND ROBIN SUB SET
SERIAL: 60050768019B02A81000000000001A1  Reserved: No  LUN SIZE: 20.0GB
HOST INTERFACE: FC
=====
Path#   Adapter/Hard Disk                State  Mode    Select  Errors
   0 *   Scsi Port6 Bus0/Disk1 Part0     OPEN  NORMAL   51     1
   1 *   Scsi Port6 Bus0/Disk1 Part0     OPEN  NORMAL   0      0
   2 *   Scsi Port7 Bus0/Disk1 Part0     OPEN  NORMAL   0      0
   3 *   Scsi Port7 Bus0/Disk1 Part0     OPEN  NORMAL   0      0
   4     Scsi Port6 Bus0/Disk1 Part0     OPEN  NORMAL   0      0
   5     Scsi Port6 Bus0/Disk1 Part0     OPEN  NORMAL   0      0
   6     Scsi Port7 Bus0/Disk1 Part0     OPEN  NORMAL   0      0
   7     Scsi Port7 Bus0/Disk1 Part0     OPEN  NORMAL   0      0

DEV#:  1  DEVICE NAME: Disk2 Part0      TYPE: 2145  POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A81000000000001A2  Reserved: No  LUN SIZE: 20.0GB
HOST INTERFACE: FC
=====
Path#   Adapter/Hard Disk                State  Mode    Select  Errors
   0     Scsi Port6 Bus0/Disk2 Part0     OPEN  NORMAL   16     0
   1     Scsi Port6 Bus0/Disk2 Part0     OPEN  NORMAL   9      0
   2     Scsi Port7 Bus0/Disk2 Part0     OPEN  NORMAL  12     0
   3     Scsi Port7 Bus0/Disk2 Part0     OPEN  NORMAL  13     0
   4 *   Scsi Port6 Bus0/Disk2 Part0     OPEN  NORMAL   0      0
   5 *   Scsi Port6 Bus0/Disk2 Part0     OPEN  NORMAL   0      0
   6 *   Scsi Port7 Bus0/Disk2 Part0     OPEN  NORMAL   0      0
   7 *   Scsi Port7 Bus0/Disk2 Part0     OPEN  NORMAL   0      0

DEV#:  2  DEVICE NAME: Disk3 Part0      TYPE: 2145  POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A81000000000001A3  Reserved: No  LUN SIZE: 20.0GB
HOST INTERFACE: FC
=====
Path#   Adapter/Hard Disk                State  Mode    Select  Errors
   0 *   Scsi Port6 Bus0/Disk3 Part0     OPEN  NORMAL   14     0
   1 *   Scsi Port6 Bus0/Disk3 Part0     OPEN  NORMAL  10     0
   2 *   Scsi Port7 Bus0/Disk3 Part0     OPEN  NORMAL  16     0
   3 *   Scsi Port7 Bus0/Disk3 Part0     OPEN  NORMAL  10     0
   4     Scsi Port6 Bus0/Disk3 Part0     OPEN  NORMAL   0      0
   5     Scsi Port6 Bus0/Disk3 Part0     OPEN  NORMAL   0      0
   6     Scsi Port7 Bus0/Disk3 Part0     OPEN  NORMAL   0      0
   7     Scsi Port7 Bus0/Disk3 Part0     OPEN  NORMAL   0      0
```

Removing additional paths

Complete the following steps to remove additional paths to an SDDDSM device:

1. Uninstall any additional hardware on the Windows Server 2008 or Windows Server 2012 server or the SAN Volume Controller. Alternatively, you can change the zoning configuration.
2. Open Computer Management window.
 - On a Windows Server 2008: Click **Start** → **Program** → **Administrative Tools** → **Computer Management**.
 - On a Windows Server 2012: Click **Server Manager** → **Tools** → **Computer Management**.
3. Click **Device Manager**.
4. Right-click **Disk drives**.
5. Click **Scan for hardware changes**.

Note: To get latest information about disks in Windows 2012 Server Core CLI, run **Get-Disk** command in the Windows PowerShell.

6. Verify that the path is removed correctly. See “Verifying that additional paths are removed correctly.”

Verifying that additional paths are removed correctly

After removing additional paths to SDDDSM devices, verify that the additional paths have been removed correctly.

Complete the following steps to verify that the additional paths are removed correctly:

1. Open the DOS prompt window.
 - On a Windows Server 2008: Click **Start** → **Program** → **Subsystem Device Driver DSM** → **Subsystem Device Driver DSM**.
 -
 - On a Windows Server 2012: Click **Start** → **Apps** → **Subsystem Device Driver DSM**.
 - On a Windows Server 2012 Server Core CLI: Go to %PROGRAMFILES%\IBM\SDDDSM.
2. Type **datapath query adapter** and press **Enter**. The output includes information about any additional adapters that were installed. In the following example, the output shows that an additional HBA is installed:

```
Active Adapters :2
Adpt#      Name      Interface  State  Mode  Select  Errors  Paths  Active
0 Scsi Port6 Bus0    FC  NORMAL  ACTIVE  1099   1     112   112
1 Scsi Port7 Bus0    FC  NORMAL  ACTIVE   813   0     112   112
```

3. Type **datapath query device** and press **Enter**. Continuing with the earlier example, the output includes information about three devices that are attached to the SCSI path:

```

DEV#: 0 DEVICE NAME: Disk1 Part0      TYPE: 2145      POLICY: ROUND ROBIN SUB SET
SERIAL: 60050768019B02A81000000000001A1  Reserved: No  LUN SIZE: 20.0GB
HOST INTERFACE: FC
=====
Path#   Adapter/Hard Disk                State  Mode    Select  Errors
 0 *   Scsi Port6 Bus0/Disk1 Part0     OPEN   NORMAL  51      1
 1 *   Scsi Port6 Bus0/Disk1 Part0     OPEN   NORMAL  0       0
 2 *   Scsi Port7 Bus0/Disk1 Part0     OPEN   NORMAL  0       0
 3 *   Scsi Port7 Bus0/Disk1 Part0     OPEN   NORMAL  0       0
 4     Scsi Port6 Bus0/Disk1 Part0     OPEN   NORMAL  0       0
 5     Scsi Port6 Bus0/Disk1 Part0     OPEN   NORMAL  0       0
 6     Scsi Port7 Bus0/Disk1 Part0     OPEN   NORMAL  0       0
 7     Scsi Port7 Bus0/Disk1 Part0     OPEN   NORMAL  0       0

DEV#: 1 DEVICE NAME: Disk2 Part0      TYPE: 2145      POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A81000000000001A2  Reserved: No  LUN SIZE: 20.0GB
HOST INTERFACE: FC
=====
Path#   Adapter/Hard Disk                State  Mode    Select  Errors
 0     Scsi Port6 Bus0/Disk2 Part0     OPEN   NORMAL  16      0
 1     Scsi Port6 Bus0/Disk2 Part0     OPEN   NORMAL  9       0
 2     Scsi Port7 Bus0/Disk2 Part0     OPEN   NORMAL  12      0
 3     Scsi Port7 Bus0/Disk2 Part0     OPEN   NORMAL  13      0
 4 *   Scsi Port6 Bus0/Disk2 Part0     OPEN   NORMAL  0       0
 5 *   Scsi Port6 Bus0/Disk2 Part0     OPEN   NORMAL  0       0
 6 *   Scsi Port7 Bus0/Disk2 Part0     OPEN   NORMAL  0       0
 7 *   Scsi Port7 Bus0/Disk2 Part0     OPEN   NORMAL  0       0

DEV#: 2 DEVICE NAME: Disk3 Part0      TYPE: 2145      POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A81000000000001A3  Reserved: No  LUN SIZE: 20.0GB
HOST INTERFACE: FC
=====
Path#   Adapter/Hard Disk                State  Mode    Select  Errors
 0 *   Scsi Port6 Bus0/Disk3 Part0     OPEN   NORMAL  14      0
 1 *   Scsi Port6 Bus0/Disk3 Part0     OPEN   NORMAL  10      0
 2 *   Scsi Port7 Bus0/Disk3 Part0     OPEN   NORMAL  16      0
 3 *   Scsi Port7 Bus0/Disk3 Part0     OPEN   NORMAL  10      0
 4     Scsi Port6 Bus0/Disk3 Part0     OPEN   NORMAL  0       0
 5     Scsi Port6 Bus0/Disk3 Part0     OPEN   NORMAL  0       0
 6     Scsi Port7 Bus0/Disk3 Part0     OPEN   NORMAL  0       0
 7     Scsi Port7 Bus0/Disk3 Part0     OPEN   NORMAL  0       0

```

- Remove four paths and run the **datapath query device** command. Continuing with the earlier example, the output shows that three devices are attached to the SCSI path. The output includes information about the paths that are not removed.

```

DEV#: 0 DEVICE NAME: Disk1 Part0 TYPE: 2145 POLICY: ROUND ROBIN SUB SET
SERIAL: 60050768019B02A810000000000001A1 Reserved: No LUN SIZE: 20.0GB
HOST INTERFACE: FC
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 Scsi Port6 Bus0/Disk1 Part0 OPEN NORMAL 23 0
1 * Scsi Port6 Bus0/Disk1 Part0 OPEN NORMAL 0 0
2 Scsi Port7 Bus0/Disk1 Part0 OPEN NORMAL 23 0
3 * Scsi Port7 Bus0/Disk1 Part0 OPEN NORMAL 0 0

DEV#: 1 DEVICE NAME: Disk2 Part0 TYPE: 2145 POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A810000000000001A2 Reserved: No LUN SIZE: 20.0GB
HOST INTERFACE: FC
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 * Scsi Port6 Bus0/Disk2 Part0 OPEN NORMAL 0 0
1 Scsi Port6 Bus0/Disk2 Part0 OPEN NORMAL 22 0
2 * Scsi Port7 Bus0/Disk2 Part0 OPEN NORMAL 0 0
3 Scsi Port7 Bus0/Disk2 Part0 OPEN NORMAL 24 0

DEV#: 2 DEVICE NAME: Disk3 Part0 TYPE: 2145 POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A810000000000001A3 Reserved: No LUN SIZE: 20.0GB
HOST INTERFACE: FC
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 Scsi Port6 Bus0/Disk3 Part0 OPEN NORMAL 25 0
1 * Scsi Port6 Bus0/Disk3 Part0 OPEN NORMAL 0 0
2 Scsi Port7 Bus0/Disk3 Part0 OPEN NORMAL 21 0
3 * Scsi Port7 Bus0/Disk3 Part0 OPEN NORMAL 0 0

```

5. Add the four paths again and run the **datapath query device** command. Continuing with the earlier example, the output shows that three devices are attached to the SCSI path. The path numbers for the existing paths do not change. The old path numbers are reassigned to the paths that you added in this step.

```

DEV#: 0 DEVICE NAME: Disk1 Part0 TYPE: 2145 POLICY: ROUND ROBIN SUB SET
SERIAL: 60050768019B02A810000000000001A1 Reserved: No LUN SIZE: 20.0GB
HOST INTERFACE: FC

```

```

=====
Path# Adapter/Hard Disk State Mode Select Errors
0 * Scsi Port6 Bus0/Disk1 Part0 OPEN NORMAL 51 1
1 * Scsi Port6 Bus0/Disk1 Part0 OPEN NORMAL 0 0
2 * Scsi Port7 Bus0/Disk1 Part0 OPEN NORMAL 0 0
3 * Scsi Port7 Bus0/Disk1 Part0 OPEN NORMAL 0 0
4 Scsi Port6 Bus0/Disk1 Part0 OPEN NORMAL 0 0
5 Scsi Port6 Bus0/Disk1 Part0 OPEN NORMAL 0 0
6 Scsi Port7 Bus0/Disk1 Part0 OPEN NORMAL 0 0
7 Scsi Port7 Bus0/Disk1 Part0 OPEN NORMAL 0 0

```

```

DEV#: 1 DEVICE NAME: Disk2 Part0 TYPE: 2145 POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A810000000000001A2 Reserved: No LUN SIZE: 20.0GB
HOST INTERFACE: FC

```

```

=====
Path# Adapter/Hard Disk State Mode Select Errors
0 Scsi Port6 Bus0/Disk2 Part0 OPEN NORMAL 16 0
1 Scsi Port6 Bus0/Disk2 Part0 OPEN NORMAL 9 0
2 Scsi Port7 Bus0/Disk2 Part0 OPEN NORMAL 12 0
3 Scsi Port7 Bus0/Disk2 Part0 OPEN NORMAL 13 0
4 * Scsi Port6 Bus0/Disk2 Part0 OPEN NORMAL 0 0
5 * Scsi Port6 Bus0/Disk2 Part0 OPEN NORMAL 0 0
6 * Scsi Port7 Bus0/Disk2 Part0 OPEN NORMAL 0 0
7 * Scsi Port7 Bus0/Disk2 Part0 OPEN NORMAL 0 0

```

```

DEV#: 2 DEVICE NAME: Disk3 Part0 TYPE: 2145 POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A810000000000001A3 Reserved: No LUN SIZE: 20.0GB
HOST INTERFACE: FC

```

```

=====
Path# Adapter/Hard Disk State Mode Select Errors
0 * Scsi Port6 Bus0/Disk3 Part0 OPEN NORMAL 14 0
1 * Scsi Port6 Bus0/Disk3 Part0 OPEN NORMAL 10 0
2 * Scsi Port7 Bus0/Disk3 Part0 OPEN NORMAL 16 0
3 * Scsi Port7 Bus0/Disk3 Part0 OPEN NORMAL 10 0
4 Scsi Port6 Bus0/Disk3 Part0 OPEN NORMAL 0 0
5 Scsi Port6 Bus0/Disk3 Part0 OPEN NORMAL 0 0
6 Scsi Port7 Bus0/Disk3 Part0 OPEN NORMAL 0 0
7 Scsi Port7 Bus0/Disk3 Part0 OPEN NORMAL 0 0

```

Uninstalling SDDDSM

Attention: If you are not planning to reinstall SDDDSM after the uninstallation, ensure that there is a single-path connection from the system to the storage device before performing a restart to avoid any potential data loss.

Complete the following steps to uninstall SDDDSM on a Windows Server 2008 or Windows Server 2012 host system:

1. Log on as the administrator user.
2. Uninstall SDDDSM.
 - On Windows Server 2008:
 - a. Click **Start** → **Settings** → **Control Panel**. The Control Panel opens.
 - b. Double-click **Add/Remove Programs**. The Add/Remove Programs window opens.
 - c. In the Add/Remove Programs window, select **Subsystem Device Driver DSM** from the currently installed programs selection list.
 - d. Click **Add/Remove**. You will be asked to confirm that you want to uninstall.
 - On Windows Server 2012:
 - a. Click **Start Screen** → **Control Panel**. The Control Panel opens.

- b. Click **Programs** → **Program and Features** → **Uninstall a program**.
- c. From the list of programs, select **Subsystem Device Driver DSM** and in the Confirm dialog box, click **OK**.
- On Windows Server 2012 Server Core CLI:
 - a. Go to %PROGRAMFILES%\IBM\SDDDSM.
 - b. Run **setup -u**.
3. Shut down the host system after the uninstallation process completes.
4. Change the zoning configuration or cable connections to ensure that there is only single-path connection from the system to the storage device.
5. Power on the host system.

SAN boot support

Use the following procedures for SAN boot support.

Remote boot support for 32-bit Windows Server 2008 or Windows Server 2012 using a QLogic HBA

Complete the following steps to install SDDDSM :

1. Configure the SAN Environment so that both Qlogic HBAs in the server can see the SAN boot device.
2. Start the server with 2 fibre-channel cables connected to both HBAs.
3. Press **Ctrl+Q** to go to Qlogic Bios Fast Utilities.
4. Select the first HBA.
5. Select **Configuration Settings**.
6. Select **Host Adapter Setting**. Enable the BIOS.
7. Press **ESC**.
8. Select **Selectable Boot Settings**.
9. Enable **Selectable Boot**.
10. Select **first (primary) boot** and press **Enter**.
11. Select **IBM storage device** and press **Enter**.
12. At the **Select LUN** prompt, select the first supported LUN, which is LUN 0.
13. Press **Esc** and select **Save Changes**.
14. Select the second HBA and repeat steps 5-13.
15. Remove the fibre-channel cable from the second HBA so that you have only a single path to first HBA.
16. Restart the Windows Server 2008 or Windows Server 2012 with the latest Service Pack CD-ROM.
17. Continue with the Windows Server 2008 or Windows Server 2012 installation process. Remember to select the SAN device that is seen by Qlogic HBA as the device to install Windows Server 2008 or Windows Server 2012. Continue with the OS installation.
18. After Windows Server 2008 or Windows Server 2012 is installed on the SAN boot device, shut down the system.
19. Unplug the fibre-channel cable from first HBA and plug the fibre-channel cable to the second HBA so that you have a single path from second HBA to the SAN device.
20. Restart the server. The system should come up in SAN Boot mode.
21. Install the latest SDDDSM and restart.

22. To add multipath support to a SAN boot device:
 - a. Shut down the server.
 - b. Plug in the fibre-channel cable to the other HBA.
 - c. Configure the SAN to have more paths to SAN boot device if needed.
23. Restart the server.

Booting from a SAN device with Windows Server 2008 or Windows Server 2012 and the SDDDSM using an Emulex HBA

Note: The **Automatic LUN Mapping** checkbox of the Emulex Configuration Settings should be selected so that both HBA ports can see all assigned LUNs.

Complete the following steps.

1. Configure the SAN Environment so that both Emulex HBAs in the server can see the SAN boot device.
2. Boot the server with 2 fibre-channel cables connected to both HBAs.
3. Press **Alt+E** to go to the Emulex BIOS Utilities.
4. Select the first HBA.
5. Select **Configure HBA Parameter Settings**.
6. Select **Option 1** to enable BIOS for this HBA.
7. Press **Page Up** to go back. Then select **Configure Boot Device**.
8. Select the first unused boot device for **Select Boot Entry** from the **List Of Saved Boot Devices**.
9. Select **01** for **Select The Two Digit Number Of The Desired Boot Device**.
10. Enter **00** for **Enter Two Digit Of Starting LUNs (hexadecimal)**.
11. Select device number **01** for **Enter Selection For Starting LUN**.
12. Select **Boot Device Via WWPn**.
13. Page up. Then select the second HBA. Repeat steps 5-12 to configure boot support for this HBA.
14. Unplug the fibre-channel cable from second HBA and plug the fibre-channel cable to the first HBA so that you have a single path from first HBA to the SAN device.
15. Restart the Windows Server 2008 or Windows Server 2012 with the latest Service Pack CD-ROM.
16. Continue with the Windows Server 2008 or Windows Server 2012 installation process.
17. After Windows Server 2008 or Windows Server 2012 is installed on the SAN Boot device, shut down the system.
18. Unplug the fibre-channel cable from the first HBA and plug in the fibre-channel cable to the second HBA so that you have a single path from second HBA to SAN device.
19. Restart the server. The system should be up in SAN boot mode.
20. Install the latest SDDDSM and restart.
21. To add multipath support to a SAN boot device:
 - a. Shut down the server.
 - b. Plug in the fibre-channel cable to the other HBA.
 - c. Configure the SAN to have more paths to the SAN boot device if needed.
22. Restart the server.

Support for Windows Server 2008 or Windows Server 2012 clustering

Configuring a Windows Server 2008 or Windows Server 2012 cluster with SDDDSM installed

The following variables are used in this procedure:

server_1

Represents the first server with two HBAs.

server_2

Represents the second server with two HBAs.

Complete the following steps to configure a Windows Server 2008 or Windows Server 2012 cluster with SDDDSM:

1. On both *server_1* and *server_2*, configure SAN devices on supported storage as shared for all HBAs.
2. Install the latest SDDDSM on *server_1*. For installation instructions, see “Installing SDDDSM” on page 341.
3. Connect fibre-channel cables from *server_1* to the supported storage, and restart *server_1*.
4. Use the **datapath query adapter** and **datapath query device** commands to verify the correct number of SAN devices and paths on *server_1*.
5. Click **Start** → **All Programs** → **Administrative Tools** → **Computer Management**. From the Computer Management window, select **Storage** and then select **Disk Management** to work with the storage devices attached to the host system.
6. Format the raw devices with NTFS and assign drive letters for all SAN devices that are going to be used as MSCS resources. Ensure that you keep track of the assigned drive letters on *server_1*.
7. Shut down *server_1*.
8. Install the latest SDDDSM on *server_2*. For installation instructions, see “Installing SDDDSM” on page 341.
9. Connect fibre-channel cables from *server_2* to the supported storage, and restart *server_2*.
10. Use the **datapath query adapter** and **datapath query device** commands to verify the correct number of SAN devices and paths on *server_2*.
11. Click **Start** → **All Programs** → **Administrative Tools** → **Computer Management**. From the Computer Management window, select **Storage** and then select **Disk Management** to work with the storage devices attached to the host system. Verify that the assigned drive letters for MSCS resources on *server_2* match the assigned drive letters on *server_1*.
12. For Windows 2008 and Windows 2012, enable the “Failover Clustering” feature and configure MSCS on *server_2*.
13. Restart *server_1*.
14. For Windows 2008 and Windows 2012, enable the “Failover Clustering” feature and configure MSCS on *server_1* as the second node of the MSCS cluster.
15. Use the **datapath query adapter** and **datapath query device** commands to verify that the correct number of LUNs and paths on *server_1* and *server_2*. (This step is optional.)

Note: You can use the **datapath query adapter** and **datapath query device** commands to show all the physical and logical volumes for the host server. The secondary server shows only the physical volumes and the logical volumes that it owns.

Removing SDDDSM in a two-node cluster environment

Use the following steps if you intend to remove the multipathing functions to a SAN device in a two-node cluster environment.

Complete the following steps to remove SDDDSM in a two-node cluster environment:

1. Move all cluster resources from node A to node B.
2. Ensure that there is a single-path connection from the system to the storage device which may include the following activities:
 - a. Disable access of second HBA to the storage device.
 - b. Change the zoning configuration to allow only one port accessed by this host.
 - c. Remove shared access to the second HBA through the IBM TotalStorage Expert V.2.1.0 Specialist.
 - d. Remove multiple SAN Volume Controller port access, if applicable.
3. Uninstall SDDDSM. See “Uninstalling SDDDSM” on page 350 for instructions.
4. Restart your system.
5. Move all cluster resources from node B to node A.
6. Complete steps 2 - 5 on node B.

About I/O timeout errors

When you are troubleshooting issues with storage performance, application hangs, and unexpected I/O timeout errors, check whether settings related to application timeout, Windows default disk I/O timeout, and adapter timeout values meet the operational requirements.

Contact technical support teams of Microsoft and other vendors for any queries related to timeout settings.

SDDDSM datapath command support

SDDDSM supports the following **datapath** commands:

- **datapath query adapter** [*<n>*] [-itype<i>iscsi/fc>]
- **datapath query device** [*<n>*] / [[*<m>* *<n>*] [-d<i>device_model>] [-i *<x>*/*<x>* *<y>*] [-s] [-c] [-a] [-tpgs] [-itype<i>iscsi/fc>]
- **datapath query adaptstats** [*<n>*]
- **datapath query devstats** [*<n>*] / [[*<m>* *<n>*] [-d<i>device_model>] [-i *<x>*/*<x>* *<y>*] [-c]
- **datapath query wwpn**
- **datapath query version**
- **datapath query portmap**
- **datapath query essmap**
- **datapath set adapter** *<n>* online/offline
- **datapath set device** *<n>* path *<m>* online/offline

- **datapath set device** <n>/(<m> <n>) policy fo/lb/sqst/lqd/lbytes/rr/rsss/df
 - fo - failover(no load balance)
 - lb - least number of I/O and adapter weight
 - sqst - shortest queue service time
 - lqd - least queue depth
 - lbytes - least bytes
 - rr - round robin
 - rrss - round robin with subset
 - df - default (uses lb)
- **datapath set performance** [on/off] [device <n>]
- **datapath clear device** <n>/(<m> <n>) count error /all
- **datapath disable/enable ports** <connection> ess <essid>
 - Supported connection formats:
 - single port = R1-Bx-Hy-Zz
 - all ports on card = R1-Bx-Hy
 - all ports on bay = R1-Bx
 - See the portmap output for the connection string and ESS serial number.
 - Examples of valid device model include:
 - 2145 - Displays all 2145 devices (SVC)
 - 2107 - Displays all 2107 devices (DS 8000)

Note: For RSSM devices, even when two or more serial-attached SCSI (SAS) HBAs are installed on the host, SDDDSM finds only a single HBA, and the output of datapath query adapter shows only one adapter.

IBM does not support RSSM. For more information, contact the IBM Support team.

For additional information about the **datapath** commands, see Chapter 10, “Using the datapath commands,” on page 371.

SDDDSM server daemon

The SDDDSM server (also referred to as sddsrv) is an integral component of SDDDSM. This component consists of a Windows application daemon that is installed in addition to the SDDDSM device driver.

Verifying if the SDDDSM server has started

After you have installed SDDDSM, complete the following steps to verify that the SDDDSM server (sddsrv) has automatically started:

1. Click **Start** → **Programs** → **Administrative Tools** → **Computer Management**.
2. Expand the Services and Applications tree.
3. Click **Services**.
4. Right-click **SDD Service**.
5. Verify that the status of the SDD Service has *Started*.

Starting the SDDDSM server manually

If the SDDDSM server did not start automatically after you performed the SDDDSM installation, complete the following steps to start sddsrv:

1. Click **Start** → **Programs** → **Administrative Tools** → **Computer Management**.

2. Expand the Services and Applications tree.
3. Click **Services**.
4. Right-click **SDD Service**.
5. Click **Start**.

Note: Alternatively, to start the SDD service, run **net start SDD_Service** on the command prompt.

Changing to a different port number for the SDDDSM server

To change to a different port number for the SDDDSM server, see the instructions in the “Changing the sddsrv or pcmsrv TCP/IP port number” on page 369.

Stopping the SDDDSM server

To stop the SDDDSM server, complete the following steps:

1. Click **Start** → **Programs** → **Administrative Tools** → **Computer Management**.
2. Expand the Services and Applications tree.
3. Click **Services**.
4. Right-click **SDD Service**.
5. Click **Stop**.

Note: Alternatively, to stop the SDD service, run **net stop SDD_Service** on the command prompt.

SDDDSM 2.4.5.0 features

The SDDDSM 2.4.5.0 release includes the following features.

Additional load balance policies in SDDDSM 2.4.5.0

Apart from the Round Robin and Optimized load balance policies, SDDDSM now includes additional load balance policies to enhance the I/O performance and to provide the user with an option to choose the policy that better suits their current SAN configuration.

Note: SVC recommends using only Active-Optimized paths for I/O as long as they are available. This is done with least number of I/O and adapter weight (lb) and Round Robin with Subset (rrss) policies. Rest of the load balance policies use all the Active-Optimized and Active-UnOptimized paths. Choose the one which suits your SAN configuration for enhanced I/O Performance. The default is set to least number of I/O and adapter weight(lb) which used to be the same as Load Balance(lb) in the earlier releases of SDDDSM.

The following are the additional load balance policies available in SDDDSM 2.4.5.0 and later releases.

Shortest Queue Service Time (sqst)

This load balance policy selects the most optimal path for the multipath device (mpdInfo) based on the pending I/O service time which is cumulative for all the outstanding I/Os together on this pathInfo. Service time for an I/O is the time taken to serve it. It is the difference of the time since MPIO has submitted the I/O to the lower driver in the device stack and the time when this I/O comes back to MPIO from storage device. **GmLBalanceShortestQueueServiceTime()** selects the new

| path with the least pending I/O service time.

| **Least Queue Depth (lqd)**

| This load balance policy selects the most optimal path for the multipath device
| (mpdInfo) based on the least queue depth which is maintained per failover group
| or per path id. Queue depth is the number of outstanding I/Os currently waiting
| in the queue to be served on this fail over group. **GmLBalanceLeastQueueDepth()**
| selects the new path with the least queue depth.

| **Least Bytes (lbytes)**

| This load balance policy selects the most optimal path for the multipath device
| (mpdInfo) based on the the least current outstanding bytes in the I/O queue for
| each path. Total bytes is considered for all the I/Os together on this device path.
| **GmLBalanceLeastBytes()** selects the new path with the least bytes.

| **Round Robin (rr)**

| This load balance policy selects the new path in a round robin manner. The last
| path used is not selected next.

| **Round Robin with Subset (rrss)**

| This load balance policy selects the most optimal path for the multipath device
| (mpdInfo) in a round robin manner within a subset of paths. The paths can be
| only from the active-optimized controller. The last path used is not selected next.

| **least number of I/O and adapter weight (lb)**

| This load balance policy selects the most optimal path for the multipath device
| (mpdInfo) based on a path with no pending I/O and that has the least
| accumulated path weight on that adapter. **GmLBalanceOptimized()** selects path with
| no I/O and the least adapter path weight.

| **Changes in datapath command**

| The **datapath** command changes for the additional load balance policies included
| in SDDDSM 2.4.5.0 are as follows.

```
| datapath set device <n>/(<m> <n>) policy fo/lb/sqst/lqd/lbytes/rr/rrss/df\n\
| fo      - failover(no load balance)\n\
| lb      - least number of I/O and adapter weight\n\
| sqst    - shortest queue service time\n\
| lqd     - least queue depth\n\
| lbytes  - least bytes\n\
| rr      - round robin\n\
| rrss    - round robin with subset\n\
| df      - default(uses lb)\n\
```

```

Administrator: Command Prompt
datapath set device <n> path <m> online/offline
datapath set device <n>/<m> <n>> policy fo/lb/sqst/lqd/lbytes/rr/rsss/df
fo - failover(no load balance)
lb - least number of I/O and adapter weight
sqst - shortest queue service time
lqd - least queue depth
lbytes - least bytes
rr - round robin
rsss - round robin with subset
df - default(uses lb)
datapath set performance [on/off] [device <n>]
datapath clear device <n>/<m> <n>> count error/all
datapath disable/enable ports <connection> ess <essid>
Supported connection format:
single port = R1-Bx-Hy-Zz
all ports on card = R1-Bx-Hy
all ports on bay = R1-Bx
refer portmap output for the connection string and ESS serial number

Examples of valid device model include:
2145 - Displays all 2145 devices <SUC>
2107 - Displays all 2107 devices <DS 8000>

```

```

Administrator: Command Prompt
D:\>datapath.exe set device 0 policy sqst
DEU#: 0 DEVICE NAME: Disk3 Part0 TYPE: 2145 POLICY: SHORTEST QUEUE SERVICE TIME
SERIAL: 60050768019B02A81000000000000233 Reserved: No LUN SIZE: 3.0GB
HOST INTERFACE: FC
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 * Scsi Port3 Bus0/Disk3 Part0 OPEN NORMAL 0 0
1 Scsi Port3 Bus0/Disk3 Part0 OPEN NORMAL 151 1
2 * Scsi Port4 Bus0/Disk3 Part0 OPEN NORMAL 0 0
3 Scsi Port4 Bus0/Disk3 Part0 OPEN NORMAL 132 1
D:\>datapath.exe set device 0 policy lqd
DEU#: 0 DEVICE NAME: Disk3 Part0 TYPE: 2145 POLICY: LEAST QUEUE DEPTH
SERIAL: 60050768019B02A81000000000000233 Reserved: No LUN SIZE: 3.0GB
HOST INTERFACE: FC
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 * Scsi Port3 Bus0/Disk3 Part0 OPEN NORMAL 0 0
1 Scsi Port3 Bus0/Disk3 Part0 OPEN NORMAL 151 1
2 * Scsi Port4 Bus0/Disk3 Part0 OPEN NORMAL 0 0
3 Scsi Port4 Bus0/Disk3 Part0 OPEN NORMAL 132 1

```

iSCSI support

SDDDSM supports iSCSI from the 2.4.5.0 release. Administrators can use iSCSI Initiator to configure iSCSI LUNs from storage and use MPIO/SDDDSM load balance I/O.

SDDDSM works with iSCSI LUNs exactly similar how it works with FC.

Changes in datapath command

```

datapath query adapter [<n>] [-itype <iscsi/fc>]\n\
datapath query device [<n>]/<m> <n>] [-d <device_model>] \n\
[-i <x>/<x><y>]] [-s] [-c] [-a] [-tpgs] [-itype <iscsi/fc>]\n\

```

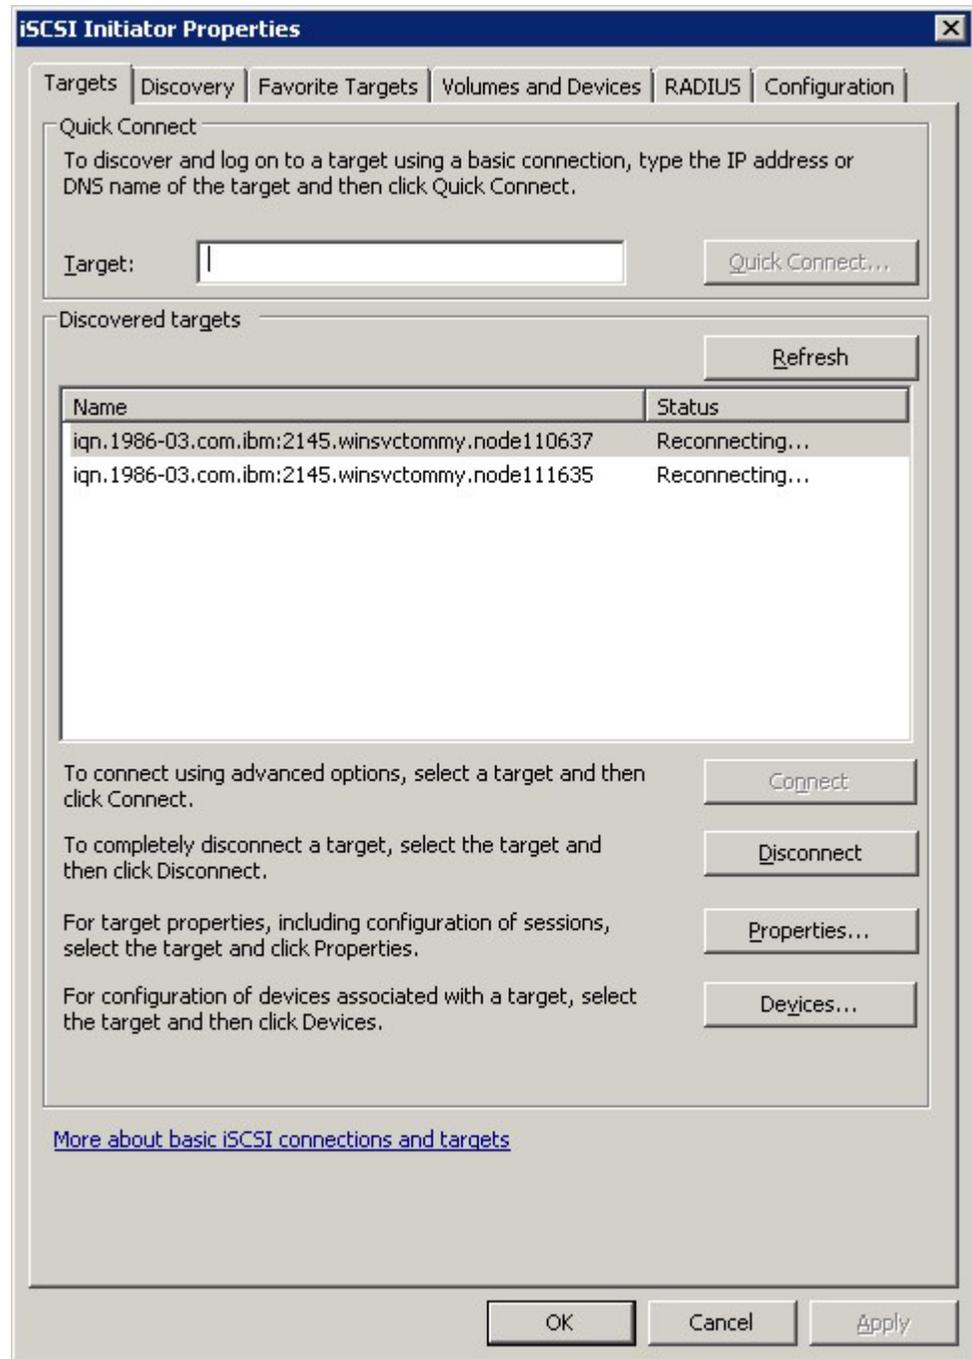
Supported platforms

The following platforms are supported.

- Windows Server 2008 R2 (64-bit)
- Windows Server 2008 R2 (32-bit)
- Windows Server 2012 (64-bit)
- Windows Server 2012 R2 (64-bit)

Mapping of LUNs

Mapping of the LUNs can be done using software iSCSI initiator which comes with operating system. iSCSI hardware initiator is not supported.



iSCSI in Hyper-V support

Microsoft supports iSCSI in Hyper-V with Windows Server 2012. The same iSCSI Initiator can be used to configure LUNs for I/O.

Changes in datapath command

```
datapath query adapter [<n>] [-itype <iscsi/fc>]\n\
datapath query device [<n>]/[<m> <n>] [-d <device_model>] \n\
[-i <x>/(<x><y>)] [-s] [-c] [-a] [-tpgs] [-itype <iscsi/fc>]\n\
```

Supported platforms

The following platforms are supported.

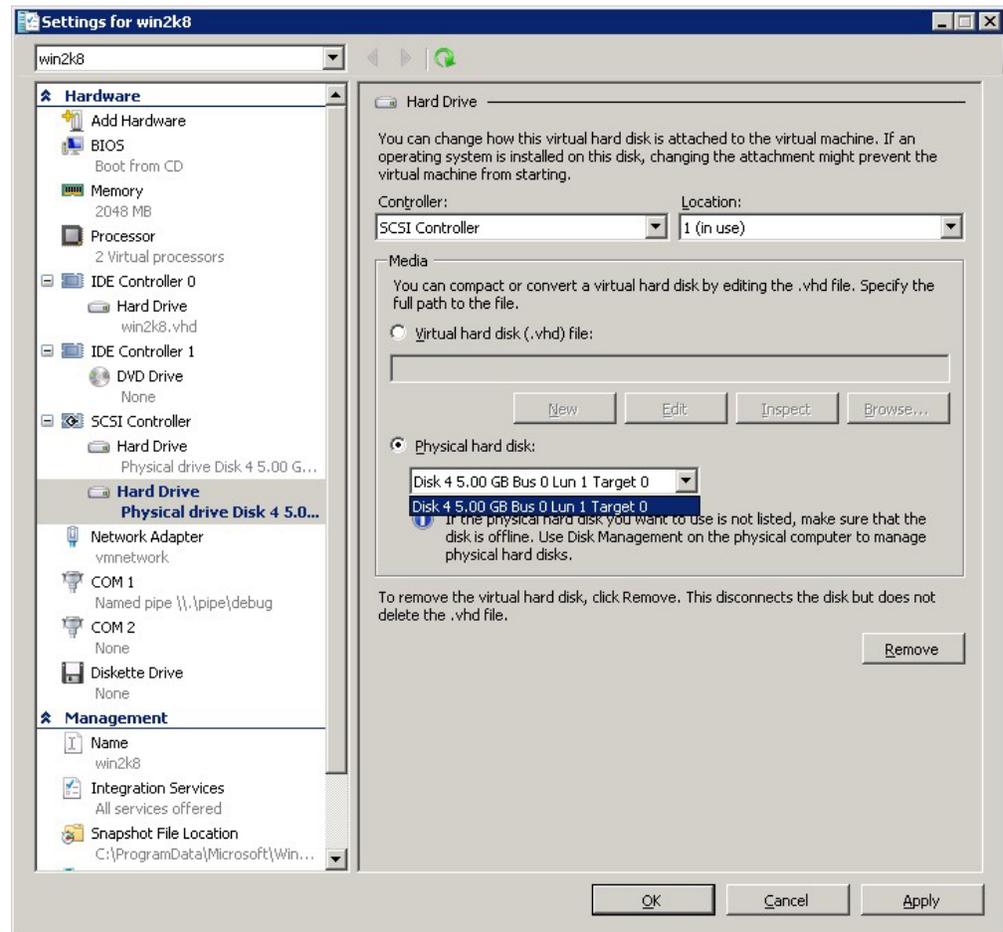
- Windows Server 2008 R2 (64-bit)
- Windows Server 2008 R2 (32-bit)
- Windows Server 2012 (64-bit)
- Windows Server 2012 R2 (64-bit)

Mapping of the LUNs

Direct mapping of the LUNs can be done using software iSCSI initiator which comes with the operating system. iSCSI hardware initiator is not supported. Pass-through LUNs come through iSCSI.

The VM location is as follows.

- VMs are installed on SAN disks.
- VMs are installed on CSVs.



FC in Hyper-V support

Microsoft supports FC in Hyper-V with MPIO from Windows Server 2008 R2 onwards. SDDDSM now supports FC in Hyper-V from the 2.4.5.0 release.

For NPIV support, refer to the HBA documentation.

Supported platforms

The following platforms are supported.

- Windows Server 2008 R2 (64-bit)
- Windows Server 2008 R2 (32-bit)
- Windows Server 2012 (64-bit)
- Windows Server 2012 R2 (32-bit)

Mapping of the LUNs

Direct mapping of the LUNs can be done using Hyper-V WWPNs from SVC. Pass-through LUNs come through the base operating system.

The VM location is as follows.

- VMs are installed on SAN disks.
- VMs are installed on CSVs and I/Os are started from them.

The following HBAs can be used.

- Emulex
- Qlogic

Note: Windows Server 2008 Hyper-V does not support FC connectivity.

Path Performance Plug-in

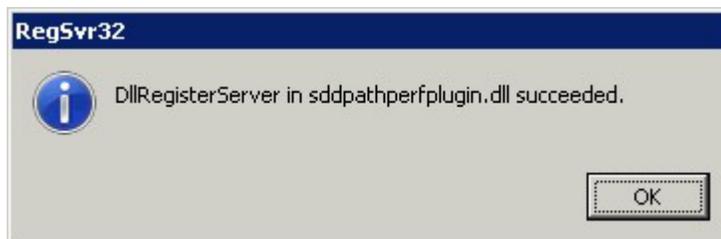
Path Performance Plug-In provides I/O statistics per individual path for any storage LUN supported by the SDDDSM driver.

SDDDSM driver is modified to gather and provide all the I/O statistics to the user mode plug-in. This plug-in comes as a DLL library and it must be registered with `regsvr32`. The Managed Object Format (MOF) file that is a part of it must be compiled with MOF compiler (`mofcomp`). Once this is done, the plug-in can interact with **PerfMon** of the Windows operating system.

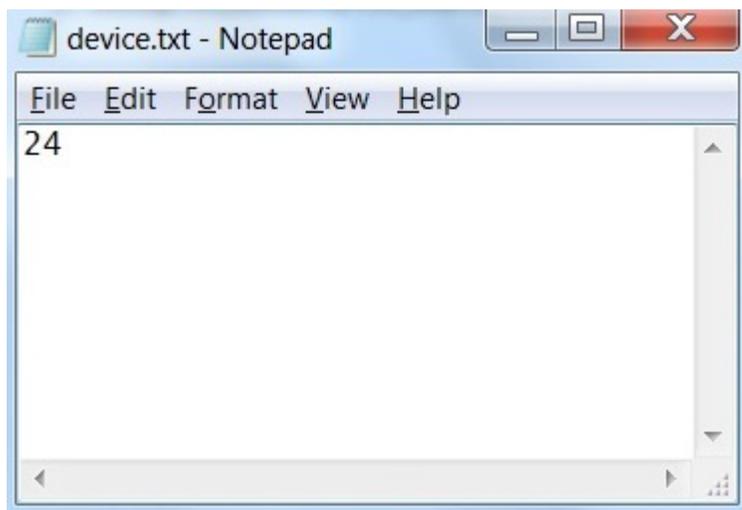
Before you can view the path performance statistics using **PerfMon**, you must complete the following steps.

1. Run `mofcomp sddpathperfplugin.MOF`.
2. Run `regsvr32 sddpathperfplugin.dll` to register the new DLL or `regsvr32 -u sddpathperf-plugin.dll` to unregister the previous DLL.

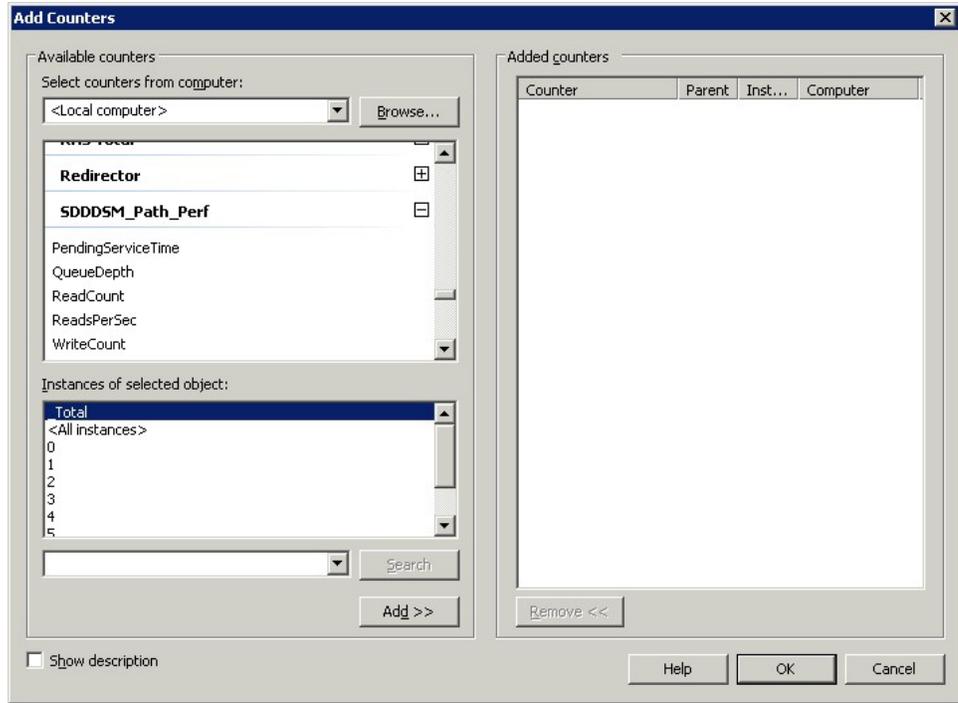
```
Administrator: Command Prompt
C:\Program Files\IBM\SDDDSM>mofcomp sddpathperfplugin.Mof
Microsoft (R) MOF Compiler Version 6.1.7600.16385
Copyright (c) Microsoft Corp. 1997-2006. All rights reserved.
Parsing MOF file: sddpathperfplugin.Mof
MOF file has been successfully parsed
Storing data in the repository...
Done!
C:\Program Files\IBM\SDDDSM>regsvr32 sddpathperfplugin.dll
C:\Program Files\IBM\SDDDSM>_
```



3. Issue the **datapath set performance on** command.
4. Issue the **datapath set device <n>** command to start generating the performance statistics for disk *n*.
5. Manually add the device number used in the preceding step to the `device.txt` file.



6. Open **PerfMon** and select **SDDDSM_Path_Perf** and its counters.



Enhancement of failover group limit per adapter

SDDDSM releases before 2.4.5.0 limited the number of failover groups (unique SCSI addresses) per adapter to 32.

This enhancement removes this limitation and users can configure any number of unique initiator-target paths through only one of the adapter. However, it is recommended to distribute I/Os between multiple adapters for better performance.

Chapter 9. Using the SDDDSM or the SDDPCM daemon

The SDD server (sddsrv) is an application program that is installed in addition to SDD.

SDDPCM server (pcmsrv) is an integrated component of SDDPCM 2.0.1.0 (or later).

SDD server daemon

The SDD server daemon (sddsrv) starts automatically after the SDD driver package is installed. The sddsrv daemon runs in the background at all times. The daemon scans to find failing paths (INVALID, CLOSE_DEAD, or DEAD) at regular intervals between two and five minutes unless otherwise indicated for a specific platform. The daemon probes idle paths that are in the CLOSE or OPEN states at regular, one-minute intervals unless otherwise indicated for a specific platform. See the appropriate chapter in this guide for the specific platform modifications to sddsrv operation.

The server daemon also provides the First Time Data Capture (FTDC) function, which periodically saves the SDD driver internal trace into a log file. Stopping the SDD server daemon loses this FTDC function, which makes debugging much harder.

Note: The sddsrv daemon is not available on NetWare host systems.

Understanding how the SDD server daemon works

The sddsrv daemon provides path reclamation and path probing.

Path reclamation

The SDD server regularly tests and recovers broken paths that have become operational. It tests invalid, close_dead, or dead paths and detects if these paths have become operational. The daemon sleeps for three-minute intervals between consecutive runs unless otherwise specified for a specific platform. If the test succeeds, sddsrv reclaims these paths and changes the states of these paths according to the following characteristics:

- If the state of the SDD vpath device is OPEN, sddsrv changes the states of INVALID and CLOSE_DEAD paths of that SDD vpath device to OPEN.
- If the state of the SDD vpath device is CLOSE, sddsrv changes the states of CLOSE_DEAD paths of that SDD vpath device to CLOSE.
- The sddsrv daemon changes the states of dead paths to OPEN.

Path probing

The SDD server regularly tests close paths and open paths that are idle to see if they are operational or have become not operational. The daemon sleeps for one-minute intervals between consecutive runs unless otherwise specified for a specific platform. If the test fails, sddsrv then changes the states of these paths according to the following characteristics:

- If the SDD vpath device is in the OPEN state and the path is not working, sddsrv changes the state of the path from OPEN to DEAD.

- If the SDD vpath device is in the CLOSE state and the path is not working, sddsrsv changes the state of the path from CLOSE to CLOSE_DEAD.
- The sddsrsv daemon will set the last path to DEAD or CLOSE_DEAD depending upon the state of the SDD vpath device.

Note: The sddsrsv daemon will not test paths that are manually placed offline.

In SDD 1.5.0.x (or earlier), sddsrsv by default was binding to a TCP/IP port and listening for incoming requests. In SDD 1.5.1.x (or later), sddsrsv does not bind to any TCP/IP port by default, but allows port binding to be dynamically enabled or disabled. For all platforms except Linux, the SDD package ships a template file of sddsrsv.conf that is named sample_sddsrsv.conf. On all UNIX platforms except Linux, the sample_sddsrsv.conf file is located in the /etc directory. On Windows platforms, the sample_sddsrsv.conf file is in the directory where SDD is installed. You must use the sample_sddsrsv.conf file to create the sddsrsv.conf file in the same directory as sample_sddsrsv.conf by copying it and naming the copied file sddsrsv.conf. You can then dynamically change port binding by modifying parameters in sddsrsv.conf.

Because the TCP/IP interface of sddsrsv is disabled by default, you cannot get sddsrsv traces from a Web browser like you could in SDD releases earlier than 1.5.1.0. Starting with SDD 1.5.1.x, the sddsrsv trace is saved in the sddsrsv log files.

The sddsrsv trace log files are wrap-around files, and the size of each file can be a maximum of 4 MB. The sddsrsv daemon also collects the SDD driver trace and puts it in log files. The daemon creates two sdd log files (sdd.log and sdd_bak.log) for the driver trace. The SDD driver trace log files are also wrap-around files, and the size of each file can be a maximum of 4 MB.

You can find sddsrsv and sdd log files in the following directories based on your host system platform:

- AIX: /var/adm/ras
- HP-UX: /var/adm/IBMsdd
- Linux: /var/log
- Solaris: /var/adm
- Windows: %windir%\system32

See Appendix A, “SDD, SDDPCM, and SDDDSM data collection for problem analysis,” on page 393 for information about reporting SDD problems.

sddsrsv and the IBM TotalStorage Expert V.2.1.0

The IBM TotalStorage Expert V.2.1.0 needs to communicate with sddsrsv through a TCP/IP socket on the port on which sddsrsv is running. The sddsrsv TCP/IP port must be enabled to listen over the network when the IBM TotalStorage Expert V.2.1.0 is collecting host volume data. Apply your corporate security rules to this port.

sddsrsv and IBM TotalStorage support for Geographically Dispersed Sites for Microsoft Cluster Service

The sddsrsv TCP/IP port must be enabled to listen over the network if you are using IBM TotalStorage Support for Geographically Dispersed Sites for Microsoft Cluster Service (MSCS). Apply your corporate security rules to this port.

SDDPCM server daemon

The SDDPCM server daemon (pcmsrv) component consists of a UNIX application daemon that is installed in addition to the SDDPCM path control module.

The pcmsrv daemon only provides the path-reclamation function for SDDPCM. It regularly tests and recovers broken paths that have become operational. By default, it tests CLOSE_FAILED paths for devices that are in the CLOSED state. It also tests OPEN_FAILED paths when a device internal healthcheck is turned off.

Note: For RSSM devices, pcmsrv will not test CLOSED_FAILED paths and always change the path state from CLOSED_FAILED to CLOSED regardless of whether the path is operational or not.

Starting with SDDPCM v2123, the pcmsrv path healthcheck adds an enhancement that automatically starts or stops the path healthcheck of a device when you issue the **pcmpath set device hc_interval** command to disable or enable a device's internal healthcheck function. However, the pcmsrv path healthcheck function does not check the path that is in good open state, which the internal path healthcheck function does. The pcmsrv daemon sleeps for one-minute intervals between consecutive runs. If the test succeeds, the pcmsrv daemon reclaims the path and changes the state of the path according to the following characteristics:

- If the state of the device is OPEN and the healthcheck function is turned off, pcmsrv changes the state of OPEN_FAILED path of that device to OPEN.
- If the state of the device is CLOSE, pcmsrv changes the state of CLOSE_FAILED path of the device to CLOSE.

The pcmsrv daemon also provides the First Time Data Capture (FTDC) function, which periodically saves the SDDPCM internal trace into a log file. When you stop the pcmsrv daemon, the FTDC function is lost, which makes debugging much harder.

The pcmsrv daemon does not bind to any TCP/IP port by default but allows port binding to be dynamically enabled or disabled. The SDDPCM package ships a template file of pcmsrv.conf that is named as sample_pcmsrv.conf. The sample_pcmsrv.conf file is located in the /etc directory. You must use the sample_pcmsrv.conf file to create the pcmsrv.conf file in the /etc directory by copying sample_pcmsrv.conf and naming the copied file pcmsrv.conf. You can then dynamically change port binding by modifying parameters in pcmsrv.conf. The trace for the pcmsrv daemon is saved in the pcmsrv log files. These are wrap-around files, and the size of each file can be a maximum of 4 MB. Trace files are located in the /var/adm/ras directory.

Starting with SDDPCM 2.6.4.x on AIX, the logging feature is enhanced. The maximum size of a pcm or pcmsrv log file is 25 MB and the maximum log file count is four. The default size of a pcm or pcmsrv log file is 4 MB and the default count is two.

sddsrv.conf file format

The sddsrv.conf file contains the following parameters:

enableport

Use this parameter to enable or disable sddsrv to bind to a TCP/IP port. The default value of this parameter is set to *false* (disabled). You can set this parameter to *true* if you want to enable the TCP/IP interface of sddsrv.

loopbackbind

If you set the `enableport` parameter to *true*, the `loopbackbind` parameter specifies whether `sddsrv` listens to any Internet address or the loopback (127.0.0.1) address. To enable `sddsrv` to listen to any Internet address, set the `loopbackbind` parameter to *false*. To enable `sddsrv` to listen only to the loopback address 127.0.0.1, the `loopbackbind` parameter must be set to *true*.

max_log_count

This parameter specifies the number of log files. You can comment out the `max_log_count` parameter to use the default value of 1. You can uncomment this parameter to change the file count to a value in the range 1 - 3. If you specify a value smaller than 1, `sddsrv` uses the default value of 1. If you specify a value greater than 3, `sddsrv` uses 3.

max_log_size

This parameter specifies the size of the log file in MB. You can comment out the `max_log_size` parameter to use the default value of 4. You can uncomment this parameter to change the file size to a value in the range 4 - 25. If you specify a value smaller than 4, `sddsrv` uses the default value of 4. If you specify a value greater than 25, `sddsrv` uses 25.

portnumber

This parameter specifies the port number that `sddsrv` binds to. The default value of this parameter is 20001. You can modify this parameter to change the port number. If the `enableport` parameter is set to *true*, you must set this parameter to a valid port number to which `sddsrv` can bind. Use a port number that is not used by any other application.

probeinterval

This parameter specifies the probe interval for `sddsrv` probing in minutes. You can leave the `probeinterval` parameter commented to use the default probe interval, which is documented in the `sample_sddsrv.conf` file. You can uncomment this parameter to change the default probe interval to a value from 0 to 65535. If you specify a value less than 0, `sddsrv` uses the default value. If you specify a value greater than 65535, `sddsrv` uses 65535.

probe_retry

This parameter specifies the number of additional retries after a SCSI inquiry by `sddsrv` fails while probing inquiries. The `probe_retry` parameter is only available for Solaris SDD. You can leave the `probe_retry` parameter commented to use the default value of 2. You can uncomment this parameter to change the number of probe retries to a value from 2 to 5. If you specify a value below 2, `sddsrv` uses the default value of 2. If you specify a value greater than 5, `sddsrv` uses 5.

You can modify these parameters while `sddsrv` is running. Your changes take effect dynamically within 30 seconds.

pcmsrv.conf file format

The `pcmsrv.conf` file contains the following parameters:

enableport

Use this parameter to enable or disable `pcmsrv` to bind to a TCP/IP port. The default value of this parameter is set to *false* (disabled). To enable the TCP/IP interface of `pcmsrv`, set this parameter to *true*.

max_log_count

This parameter specifies the number of log files. You can comment out the

`max_log_count` parameter to use the default value of 1. You can uncomment this parameter to change the file count to a value in the range 1 - 3. If you specify a value smaller than 1, `pcmsrv` uses the default value of 1. If you specify a value greater than 3, `pcmsrv` uses 3.

max_log_size

This parameter specifies the size of the log file in MB. You can comment out the `max_log_size` parameter to use the default value of 4. You can uncomment this parameter to change the file size to a value in the range 4 - 25. If you specify a value smaller than 4, `pcmsrv` uses the default value of 4. If you specify a value greater than 25, `pcmsrv` uses 25.

loopbackbind

If you set the `enableport` parameter to *true*, the `loopbackbind` parameter specifies whether `pcmsrv` listens to any Internet address or the loopback (127.0.0.1) address. To enable `pcmsrv` to listen to any Internet address, set the `loopbackbind` parameter to *false*. To enable `pcmsrv` to listen only to the loopback address 127.0.0.1, set the `loopbackbind` parameter to *true*.

portnumber

This parameter specifies the port number that `pcmsrv` binds to. The default value of this parameter is 20001. You can modify this parameter to change the port number. If the `enableport` parameter is set to *true*, set this parameter to a valid port number to which `pcmsrv` can bind. Use a port number that is not used by any other application.

You can modify these parameters while `pcmsrv` is running. Your changes take effect dynamically.

Enabling or disabling the `sddsrv` or `pcmsrv` TCP/IP port

By default, `sddsrv` and `pcmsrv` do not bind to any TCP/IP port because the `enableport` parameter defaults to a value of *false*. However, you can enable or disable port binding by changing the `enableport` parameter in the `sddsrv.conf/pcmsrv.conf` file.

`enableport = true` enables `sddsrv` or `pcmsrv` to bind to a TCP/IP port.

`enableport = false` disables `sddsrv` or `pcmsrv` from binding to a TCP/IP port.

Changing the `sddsrv` or `pcmsrv` TCP/IP port number

You can modify the `portnumber` parameter in the configuration file to change the port number to which `sddsrv` or `pcmsrv` can bind. Use a port number that is not used by any other application. If the `enableport` parameter is set to *true*, only `sddsrv` or `pcmsrv` will bind to the port number specified. The default value of this parameter is 20001.

Disabling the probing function in `sddsrv`

By default, `sddsrv` probing is enabled. If you want to disable probing in `sddsrv`, uncomment the `probeinterval` parameter and set it to 0, as shown in the following example:

```
probeinterval = 0
```

If you disable sddsrsv probing, sddsrsv is not able to mark bad ideal paths to the DEAD or CLOSED_DEAD state, which might lengthen cluster failover. Also, it takes more time to open SDD devices with bad paths.

Changing the probing interval in sddsrsv

By default, sddsrsv probing is enabled with the default value as documented in the sample_sddsrsv.conf file. To change the sddsrsv probing interval, uncomment the probeinterval parameter and set it to a valid value. The following example shows how to set the interval to 5 minutes.

```
probeinterval = 5
```

Note: pcmsrv and sddsrsv in SDDDSM do not have a probing function.

Chapter 10. Using the datapath commands

SDD provides commands that you can use to complete the following tasks:

- Display the status of adapters that are used to access managed devices.
- Display the status of devices that the device driver manages.
- Dynamically set the status of paths or adapters to *online* or *offline*.
- Dynamically remove paths or adapters.
- Open an invalid or close_dead path.
- Change the path selection algorithm policy of a device.
- Run the essutil Product Engineering tool.
- Dynamically clear the select counter or error counter

This chapter includes descriptions of these commands. Table 27 provides an alphabetical list of these commands, a brief description, and where to go in this chapter for more information.

Table 27. Commands

Command	Description	Page
datapath clear device count	Dynamically clears the select counter or error counter.	“datapath clear device count” on page 372
datapath disable ports	Places paths connected to certain ports offline.	“datapath disable ports” on page 373
datapath enable ports	Places paths connected to certain ports online.	“datapath enable ports” on page 374
datapath open device path	Dynamically opens a path that is in an invalid or close_dead state.	“datapath open device path” on page 374
datapath query adapter	Displays information about adapters.	“datapath query adapter” on page 376
datapath query adaptstats	Displays performance information for all SCSI and FCS adapters that are attached to SDD devices.	“datapath query adaptstats” on page 377
datapath query device	Displays information about devices.	“datapath query device” on page 378
datapath query devstats	Displays performance information for a single SDD vpath device or all SDD vpath devices.	“datapath query devstats” on page 381
datapath query essmap	Displays each SDD vpath device, path, location, and attributes.	“datapath query essmap” on page 383
datapath query portmap	Displays the status of the logic paths that are managed by SDD between the host and the storage ports.	“datapath query portmap” on page 384

Table 27. Commands (continued)

Command	Description	Page
datapath query version	Displays the version of SDD that is installed.	“datapath query version” on page 385
datapath query wwpn	Displays the World Wide Port Name (WWPN) of the host fibre-channel adapters.	“datapath query wwpn” on page 386
datapath remove adapter	Dynamically removes an adapter.	“datapath remove adapter” on page 386
datapath remove device path	Dynamically removes a path of an SDD vpath device.	“datapath remove device path” on page 387
datapath set adapter	Sets all device paths that are attached to an adapter to online or offline.	“datapath set adapter” on page 389
datapath set device policy	Dynamically changes the path-selection policy of a single or multiple SDD vpath devices.	“datapath set device policy” on page 389
datapath set device path	Sets the path of an SDD vpath device to online or offline.	“datapath set device path” on page 390
datapath set performance	Initiates or terminates the performance statistics collection in the SDDDSM driver.	“datapath set performance” on page 391
datapath set qdepth	Dynamically enables or disables queue depth of an SDD vpath device.	“datapath set qdepth” on page 391

datapath clear device count

The **datapath clear device count** command dynamically sets the select counter or error counter to zero.

Syntax

```

▶▶ datapath clear device number 1 device number 2 count error

```

Parameters

device number 1 <*device number 2*>

When two device numbers are entered, this command applies to all the devices whose index numbers fit within the range of these two device index numbers.

error

Clears only the error counter of the SDD vpath device or range of devices specified.

all

Clears both the select counter and the error counter of the SDD vpath device or devices in the specified range.

Examples

If you have a nonzero select counter or error counter and enter the **datapath query device** command, the following output is displayed:

```
DEV#: 0 DEVICE NAME: vpath0 TYPE: 2145 POLICY: Optimized
SERIAL: 600507680181006B20000000000000D1
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 fscsi0/hdisk15 CLOSE NORMAL 53020 47
1 fscsi0/hdisk20 CLOSE NORMAL 0 0
2 fscsil/hdisk55 CLOSE NORMAL 365742 0
3 fscsil/hdisk60 CLOSE NORMAL 0 0
```

If you enter the **datapath clear device 0 count all** command and then enter the **datapath query device** command, the following output is displayed:

```
DEV#: 0 DEVICE NAME: vpath0 TYPE: 2145 POLICY: Optimized
SERIAL: 600507680181006B20000000000000D1
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 fscsi0/hdisk15 CLOSE NORMAL 0 0
1 fscsi0/hdisk20 CLOSE NORMAL 0 0
2 fscsil/hdisk55 CLOSE NORMAL 0 0
3 fscsil/hdisk60 CLOSE NORMAL 0 0
```

datapath disable ports

The **datapath disable ports** command sets SDD vpath device paths offline for the specified disk storage system location code.

Note: This command is not supported for Novell host systems.

Syntax

```
▶▶—datapath disable ports—<connection>—ess <essid>—————▶▶
```

Parameters

connection

The connection code must be in one of the following formats:

- Single port = R1-Bx-Hy-Zz
- All ports on card = R1-Bx-Hy
- All ports on bay = R1-Bx

Use the output of the **datapath query essmap** command to determine the connection code.

essid

The disk storage system serial number, given by the output of the **datapath query portmap** command.

Examples

If you enter the **datapath disable ports R1-B1-H3 ess 12028** command and then enter the **datapath query device** command, the following output is displayed:

```
DEV#: 0 DEVICE NAME: vpath0 TYPE: 2105E20 POLICY: Optimized
SERIAL: 20112028
=====
Path# Adapter/Path Name State Mode Select Errors
```

0	fscsi0/hdisk2	DEAD	OFFLINE	6	0
1	fscsi0/hdisk4	OPEN	NORMAL	9	0
2	fscsi1/hdisk6	DEAD	OFFLINE	11	0
3	fscsi1/hdisk8	OPEN	NORMAL	9	0

datapath enable ports

The **datapath enable ports** command sets SDD vpath device paths online for the specified disk storage system location code.

Note: This command is not supported for Novell host systems.

Syntax

▶▶—datapath enable ports—*connection*—*ess essid*—————▶▶

Parameters

connection

The connection code must be in one of the following formats:

- Single port = R1-Bx-Hy-Zz
- All ports on card = R1-Bx-Hy
- All ports on bay = R1-Bx

Use the output of the **datapath essmap** command to determine the connection code.

essid

The disk storage system serial number, given by the output of the **datapath query portmap** command.

Examples

If you enter the **datapath enable ports R1-B1-H3 ess 12028** command and then enter the **datapath query device command**, the following output is displayed:

```
DEV#: 0 DEVICE NAME: vpath0 TYPE: 2105E20 POLICY: Optimized
SERIAL: 20112028
=====
Path#      Adapter/Path Name      State  Mode      Select  Errors
0          fscsi0/hdisk2         OPEN   NORMAL    6       0
1          fscsi0/hdisk4         OPEN   NORMAL    9       0
2          fscsi1/hdisk6         OPEN   NORMAL    11      0
3          fscsi1/hdisk8         OPEN   NORMAL    9       0
```

datapath open device path

The **datapath open device path** command dynamically opens a path that is in the INVALID or CLOSE_DEAD state. You can use this command even when the I/O is actively running.

Note: This command is supported for AIX, HP, Linux, and Sun host systems.

Syntax

► `datapath open device—device number—path—path number` ◄

Parameters

device number

The device number refers to the device *index* number as displayed by the **datapath query device** command.

path number

The path number that you want to change, as displayed by the **datapath query device** command.

Examples

If you enter the **datapath query device 8** command, the following output is displayed:

```
DEV#: 8 DEVICE NAME: vpath9 TYPE: 2105E20 POLICY: Optimized
SERIAL: 20112028
=====
Path#    Adapter/Hard Disk    State    Mode    Select    Errors
  0      fscsi1/hdisk18      OPEN     NORMAL    557       0
  1      fscsi1/hdisk26      OPEN     NORMAL    568       0
  2      fscsi0/hdisk34      INVALID  NORMAL     0         0
  3      fscsi0/hdisk42      INVALID  NORMAL     0         0
```

Note that the current state of path 2 is INVALID.

If you enter the **datapath open device 8 path 2** command, the following output is displayed:

```
Success: device 8 path 2 opened

DEV#: 8 DEVICE NAME: vpath9 TYPE: 2105E20 POLICY: Optimized
SERIAL: 20112028
=====
Path#    Adapter/Hard Disk    State    Mode    Select    Errors
  0      fscsi1/hdisk18      OPEN     NORMAL    557       0
  1      fscsi1/hdisk26      OPEN     NORMAL    568       0
  2      fscsi0/hdisk34      OPEN     NORMAL     0         0
  3      fscsi0/hdisk42      INVALID  NORMAL     0         0
```

After you issue the **datapath open device 8 path 2** command, the state of path 2 becomes *open*.

The terms used in the output are defined as follows:

Dev# The number of this device.

Device name

The name of this device.

Type The device product ID from inquiry data.

Policy The current path-selection policy selected for the device. See “datapath set device policy” on page 389 for a list of valid policies.

Serial The logical unit number (LUN) for this device.

Path# The path number displayed by the **datapath query device** command.

Adapter

The name of the adapter to which the path is attached.

Hard Disk

The name of the logical device to which the path is bound.

State The condition of the named device:

Open Path is in use.

Close Path is not being used.

Close_Dead

Path is broken and is not being used.

Dead Path is no longer being used.

Invalid

The path failed to open.

Mode The mode of the named path, which is either Normal or Offline.

Select The number of times that this path was selected for input and output.

Errors The number of input errors and output errors that are on this path.

datapath query adapter

The **datapath query adapter** command displays information about a single adapter or all adapters.

Syntax

►►—datapath query adapter—*n itype*—◀◀

Parameters

n adapter number

The index number for the adapter for which you want information displayed. If you do not enter an adapter index number, information about all adapters is displayed.

itype

Interface type

Examples

If you enter the **datapath query adapter -itype fc** command, the following output is displayed:

```
Active FC Adapters :2
Adpt# Name Interface State Mode Select Errors Paths Active
0 Scsi Port6 Bus0 FC NORMAL ACTIVE 3304 2 60 60
1 Scsi Port5 Bus0 FC NORMAL ACTIVE 3777 2 60 60
```

```
datapath.exe query adapter -itype iscsi
```

```
No active ISCSI adapter found.
```

The terms used in the output are defined as follows:

Adpt #

The number of the adapter defined by SDD.

Adapter Name

The name of the adapter.

Interface

The type of the interface.

State The condition of the named adapter. It can be either:

Normal

Adapter is in use.

Degraded

One or more paths attached to the adapter are not functioning.

Failed All paths attached to the adapter are no longer operational.

Mode The mode of the named adapter, which is either Active or Offline.

Select The number of times this adapter was selected for input or output.

Errors The number of errors on all paths that are attached to this adapter.

Paths The number of paths that are attached to this adapter.

Note: In the Windows NT host system, this is the number of physical and logical devices that are attached to this adapter.

Active The number of functional paths that are attached to this adapter. The number of functional paths is equal to the number of paths attached to this adapter minus any that are identified as failed or offline.

datapath query adaptstats

The **datapath query adaptstats** command displays performance information for all SCSI and fibre-channel adapters that are attached to SDD devices. If you do not enter an adapter number, information about all adapters is displayed.

Syntax

►►—datapath query adaptstats—*adapter number*—◀◀

Parameters

adapter number

The index number for the adapter for which you want information displayed.

If you do not enter an adapter index number, information about all adapters is displayed.

Examples

If you enter the **datapath query adaptstats 0** command, the following output is displayed:

```

Adapter #: 0
=====
                Total Read  Total Write  Active Read  Active Write  Maximum
I/O:                1442    41295166         0           2         75
SECTOR:             156209    750217654         0          32        2098

/*-----*/

```

The terms used in the output are defined as follows:

Total Read

- I/O: total number of completed read requests
- SECTOR: total number of sectors that have been read

Total Write

- I/O: total number of completed write requests
- SECTOR: total number of sectors that have been written

Active Read

- I/O: total number of read requests in process
- SECTOR: total number of sectors to read in process

Active Write

- I/O: total number of write requests in process
- SECTOR: total number of sectors to write in process

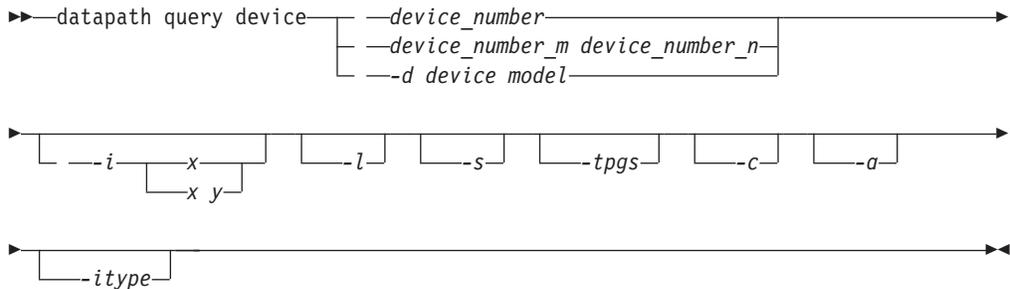
Maximum

- I/O: the maximum number of queued I/O requests
- SECTOR: the maximum number of queued sectors to *read* or *write*

datapath query device

The **datapath query device** command displays information about a single device, a range of devices, a specific model of supported storage devices, or all devices. If you do not enter a device number, information about all devices is displayed. You cannot use the option to specify a device model when you query with device numbers.

Syntax



Parameters

device_number

The device index number that is displayed by the **datapath query device** command, rather than the SDD device number.

device_number_m device_number_n

The option that you can use to provide a range of device index numbers.

-d device model

The device model that you want to display. The option to specify a device model is supported on all platforms except Novell.

Examples of valid device models include the following models:

2105 All 2105 models (ESS).

2105F All 2105 F models (ESS).

2105 800

All 2105 800 models (ESS).

2145 All 2145 models (SAN Volume Controller).

2107 All DS8000 models.

1750 All DS6000 models.

- i Repeats the command every *x* seconds for *y* times. If *y* is not specified, the command repeats every *x* seconds indefinitely.
- l Marks the nonpreferred paths with an asterisk, displays the LUN identifier, and for AIX only, displays the qdepth_enable value.
- s Queries the SCSI address of the device. This option is available for both SDD 1.6.1.x (or later) and SDDDSM 2.1.1.x (or later) for Windows platforms.

-tpgs

Shows state of all the available paths based on SCSI TPGS standards. Here is a sample output:

Path#	Adapter/Hard Disk	State	Mode	Select	Errors
4	Scsi Port6 Bus0/Disk7 Part0	Standby	OFFLINE	544	0
5	Scsi Port6 Bus0/Disk7 Part0	Active/Optimized	NORMAL	527	0
6 *	Scsi Port6 Bus0/Disk7 Part0	Active/Unoptimized	NORMAL	5	0
7 *	Scsi Port6 Bus0/Disk7 Part0	Active/Unoptimized	NORMAL	2	1

- Paths from the preferred controllers are shown as "Active Optimized".
- Paths from the non-preferred controllers are shown as "Active Unoptimized".
- Paths that are in dead state are shown as "Standby".
- To get the unavailable path states, use the option "-a".

Note: This parameter is available with SDDDSM 2.4.4.0 and later.

-c Does not show the removed device.

Note: This parameter is available with SDDDSM 2.4.4.0 and later.

-a Shows state of all the available and unavailable paths. Here is a sample output:

Path#	Adapter/Hard Disk	State	Mode	Select	Errors
0	Scsi Port5 Bus0/Disk7 Part0	CLOSE	OFFLINE	323	0
1	Scsi Port5 Bus0/Disk7 Part0	CLOSE	NORMAL	335	0
2 *	Scsi Port5 Bus0/Disk7 Part0	CLOSE	NORMAL	9	0
3 *	Scsi Port5 Bus0/Disk7 Part0	CLOSE	NORMAL	8	0
4	Scsi Port6 Bus0/Disk7 Part0	DEAD	OFFLINE	544	0
5	Scsi Port6 Bus0/Disk7 Part0	OPEN	NORMAL	527	0
6 *	Scsi Port6 Bus0/Disk7 Part0	OPEN	NORMAL	5	0
7 *	Scsi Port6 Bus0/Disk7 Part0	OPEN	NORMAL	2	1

Note: This parameter is available with SDDDSM 2.4.4.0 and later.

-i type

Specifies the interface type. The allowed values are iscsi and fc.

Note: This parameter is available with SDDDSM 2.4.5.0 and later.

Examples

If you enter the **datapath query device 0** command, the following output is displayed:

```
DEV#: 0 DEVICE NAME: vpath0 TYPE: 2145 POLICY: Optimized
SERIAL: 6005076801818008C000000000000065
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 fscsil/hdisk72 OPEN NORMAL 0 0
1 fscsi0/hdisk22 OPEN NORMAL 5571118 0
2 fscsi0/hdisk32 OPEN NORMAL 0 0
3 fscsil/hdisk62 OPEN NORMAL 5668419 0
```

If you enter the **datapath query device 0 -1** command for a device type that has preferred and non-preferred paths, the following output is displayed:

```
DEV#: 0 DEVICE NAME: vpath0 TYPE: 2145 POLICY: Optimized
SERIAL: 6005076801818008C000000000000065
LUN IDENTIFIER: 6005076801818008C000000000000065
=====
Path# Adapter/Hard Disk State Mode Select Errors
0* fscsil/hdisk72 OPEN NORMAL 0 0
1 fscsi0/hdisk22 OPEN NORMAL 5571118 0
2* fscsi0/hdisk32 OPEN NORMAL 0 0
3 fscsil/hdisk62 OPEN NORMAL 5668419 0
```

If you enter the **datapath query device -itype fc** command, the following output is displayed:

Total FC Devices : 15

```
DEV#: 0 DEVICE NAME: Disk2 Part0 TYPE: 2145 POLICY: LEAST I/O AND WEIGHT
SERIAL: 60050768019B02A810000000000001FA Reserved: No LUN SIZE: 5.0GB
HOST INTERFACE: FC
=====
Path# Adapter/Hard Disk State Mode Select Errors
0 * Scsi Port6 Bus0/Disk2 Part0 OPEN NORMAL 7 1
1 * Scsi Port6 Bus0/Disk2 Part0 OPEN NORMAL 6 1
2 * Scsi Port5 Bus0/Disk2 Part0 OPEN NORMAL 7 1
3 * Scsi Port5 Bus0/Disk2 Part0 OPEN NORMAL 8 1
4 Scsi Port5 Bus0/Disk2 Part0 OPEN NORMAL 173 0
5 Scsi Port5 Bus0/Disk2 Part0 OPEN NORMAL 96 0
6 Scsi Port6 Bus0/Disk2 Part0 OPEN NORMAL 81 0
7 Scsi Port6 Bus0/Disk2 Part0 OPEN NORMAL 90 0
```

Note:

1. Usually, the *device number* and the *device index* number are the same. However, if the devices are configured out of order, the two numbers are not always consistent. To find the corresponding index number for a specific device, always run the **datapath query device** command first.
2. For SDD 1.4.0.0 (or later), the location of the policy and serial number are swapped.

The terms that are used in the output are defined as follows:

Dev# The number of this device that is defined by SDD.

Name The name of this device that is defined by SDD.

Type The device product ID from inquiry data.

Policy The current path selection policy that is selected for the device. See “datapath set device policy” on page 389 for a list of valid policies.

Serial The LUN for this device.

Path# The path number.

Adapter

The name of the adapter to which the path is attached.

Hard disk

The name of the logical device to which the path is bound.

State The condition of the named device:

Open Path is in use.

Close Path is not being used.

Close_Dead

Path is broken and not being used.

Dead Path is no longer being used. It was either removed by SDD due to errors or manually removed using the **datapath set device M path N offline** or **datapath set adapter N offline** command.

Invalid

The path failed to open.

Mode The mode of the named path. The mode can be either Normal or Offline.

Select The number of times this path was selected for input or output.

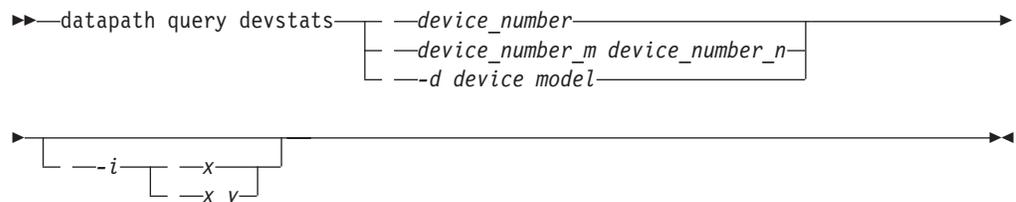
Errors The number of input and output errors on a path that is attached to this device.

datapath query devstats

The **datapath query devstats** command displays performance information for a single SDD device, a range of devices, a specific model of supported storage devices, or all SDD devices. If you do not enter a device number, information about all devices is displayed.

The option to specify a device model cannot be used when you query with device numbers.

Syntax



Parameters

device number

The device index number that is displayed by the **datapath query device** command, rather than the SDD device number.

device_number_m device_number_n

The option that you can use to provide a range of device index numbers.

-d device model

The device model that you want to display.

Note: The *-d device model* option is supported on AIX only.

Examples of valid device models include the following models:

2105 All 2105 models (ESS).

2105F All 2105 F models (ESS).

2105800
All 2105 800 models (ESS).

2145 All 2145 models (SAN Volume Controller).

2107 All DS8000 models.

1750 All DS 6000 models.

Note: The option to specify a device model is supported on all platforms except Novell.

-i Repeats the command every *x* seconds for *y* times. If *y* is not specified, the command will repeat every *x* seconds indefinitely.

Examples

If you enter the **datapath query devstats 0** command, the following output is displayed:

```
Device #: 0
=====
                Total Read  Total Write  Active Read  Active Write  Maximum
I/O:                387      24502563           0           0           62
SECTOR:              9738      448308668           0           0          2098

Transfer Size:      <= 512      <= 4k      <= 16K      <= 64K      > 64K
                   4355850      1024164      19121140      1665          130

/*-----*/
```

The terms used in the output are defined as follows:

Total Read

- I/O: total number of completed read requests
- SECTOR: total number of sectors that have been read

Total Write

- I/O: total number of completed write requests
- SECTOR: total number of sectors that have been written

Active Read

- I/O: total number of read requests in process
- SECTOR: total number of sectors to read in process

Active Write

- I/O: total number of write requests in process
- SECTOR: total number of sectors to write in process

Maximum

- I/O: the maximum number of queued I/O requests

- **SECTOR**: the maximum number of queued sectors to Read or Write

Transfer size

- <= 512: the number of I/O requests received, whose transfer size is 512 bytes or less
- <= 4k: the number of I/O requests received, whose transfer size is 4 KB or less (where KB equals 1024 bytes)
- <= 16K: the number of I/O requests received, whose transfer size is 16 KB or less (where KB equals 1024 bytes)
- <= 64K: the number of I/O requests received, whose transfer size is 64 KB or less (where KB equals 1024 bytes)
- > 64K: the number of I/O requests received, whose transfer size is greater than 64 KB (where KB equals 1024 bytes)

datapath query essmap

The **datapath query essmap** command displays data collected from the `sddfcmap` program invoked by `cfgmgr`. `sddfcmap` issues SCSI commands directly to the disks (inquiry, read capacity, and log sense) to collect the information displayed.

Note:

1. This command is not supported for Novell host systems.
2. This command is supported for disk storage systems only.
3. If the output of **datapath query essmap** does not display the correct volume group name on an HP-UX host system, complete the following steps:
 - a. Find the process ID number (PID) of `sddsrv`.
 - b. Enter `kill -SIGUSR1` (the PID of `sddsrv`).
 - c. Wait for a little while and then run the **datapath query essmap** command again.

Syntax

▶▶—datapath query essmap—▶▶

Examples

If you enter the **datapath query essmap** command, the following output is displayed:

Disk	Path	P	Location	adapter	LUN	SN	Type	Size	LSS	Vol	Rank	C/A	S ...
vpath20	hdisk1		30-60-01[FC]	fscsi1	13AAAKA1200	IBM	1750-500	1.1	18	0	0000	01	Y ...
vpath20	hdisk720	*	30-60-01[FC]	fscsi1	13AAAKA1200	IBM	1750-500	1.1	18	0	0000	01	Y ...
vpath20	hdisk848		20-60-01[FC]	fscsi0	13AAAKA1200	IBM	1750-500	1.1	18	0	0000	01	Y ...
vpath20	hdisk976	*	20-60-01[FC]	fscsi0	13AAAKA1200	IBM	1750-500	1.1	18	0	0000	01	Y ...

The terms used in the output are defined as follows:

Disk The logical device name assigned by the host.

Path The logical path name of an SDD vpath device.

P Indicates whether the logical path is a preferred path or nonpreferred path. "*" indicates that it is a nonpreferred path. This field applies only to 1750 devices.

Location

The physical location code of the host adapter through which the LUN is accessed.

Adapter

The logical adapter name that is assigned by the host LUN.

LUN SN

The unique serial number for each LUN within the disk storage system.

Type The device and model.

Size The configured capacity of the LUN.

LSS The logical subsystem where the LUN resides. (Beginning with 1.6.3.0, the value displayed is changed from decimal to hexadecimal.)

Vol The volume number within the disk storage system.

Rank The unique identifier for each RAID within the disk storage system.

C/A The cluster and adapter accessing the array.

S Indicates that the device is shared by two and more disk storage system ports. Valid values are *yes* or *no*.

Connection

The physical location code of disk storage system adapter through which the LUN is accessed.

Port The disk storage system port through which the LUN is accessed.

RaidMode

The disk RAID mode.

datapath query portmap

The **datapath query portmap** command displays the status of the logic paths that are managed by SDD between the host and the storage ports.

Note:

1. This command is not supported for Novell host systems.
2. This command is supported for disk storage systems only.
3. For Solaris, only the state of slice c of a lun is displayed. The states of other slices of the same lun might be in different states.

Syntax

▶▶—datapath query portmap—▶▶

Examples

If you enter the **datapath query portmap** command, the following output is displayed:

```

          BAY-1(B1)          BAY-2(B2)          BAY-3(B3)          BAY-4(B4)
    ESSID  DISK    H1  H2  H3  H4    H1  H2  H3  H4    H1  H2  H3  H4    H1  H2  H3  H4
          ABCD ABCD ABCD ABCD    ABCD ABCD ABCD ABCD    ABCD ABCD ABCD ABCD    ABCD ABCD ABCD ABCD
          BAY-5(B5)          BAY-6(B6)          BAY-7(B7)          BAY-8(B8)
    13AAKA vpath20  H1  H2  H3  H4    H1  H2  H3  H4    H1  H2  H3  H4    H1  H2  H3  H4
    13AAKA vpath21  ABCD ABCD ABCD ABCD    ABCD ABCD ABCD ABCD    ABCD ABCD ABCD ABCD    ABCD ABCD ABCD ABCD
    Y--- -----
    Y = online/open          y = (alternate path) online/open
  
```

O = online/closed o = (alternate path) online/closed
N = offline n = (alternate path) offline
- = path not configured
PD = path down

The terms used in the output are defined as follows:

- Y** The port is online and open. At least one path attached to this port is functional.
- y** Paths connected to this port are nonpreferred paths. The port is online and open. At least one path attached to this port is functional.
- O** The port is online and closed. At least one path state and mode is closed and online.
- o** Paths connected to this port are nonpreferred paths. The port is online and closed. At least one path state and mode is closed and online.
- N** The port is offline. All paths attached to this port are offline.
- n** Paths connected to this port are nonpreferred paths. The port is offline. All paths attached to this port are offline.
- The path is not configured.
- PD** The path is down. It is either not functional or has been placed offline.

Note: The following fields apply only to 1750 devices:

- y
- o
- n

The serial number of ESS devices is five digits, whereas the serial number of DS6000 and DS8000 devices is seven digits.

datapath query version

The **datapath query version** command displays the version of the SDD that is installed.

Syntax

▶▶—datapath query version—▶▶

Parameters

None

Examples

If you enter the **datapath query version** command, the following output is displayed:

```
[root@abc]> datapath query version  
IBM SDD Version 1.6.1.0 (devices.sdd.52.rte)
```

datapath query wwpn

The **datapath query wwpn** command displays the worldwide port name (WWPN) of the host fibre-channel adapters.

Note:

1. This command is available for AIX, Linux, HP, and Solaris host systems.
2. For Solaris, this option is supported on Solaris 8, 9, and 10. It requires that SNIA HBA libraries are installed in the Solaris system. The SNIA HBA library consists of the following two components:
 - Common SNIA HBA library
 - Vendor-specific SNIA HBA library

The common SNIA HBA library is included in the Solaris 10 installation, but users need to download the common SNIA HBA library from SUN for Solaris 8 and 9. The vendor-specific SNIA HBA library is provided by each HBA vendor.

3. Because this command is only for fibre-channel adapters, it does not display the WWPN of serial-attached SCSI (SAS) HBAs that are used to connect to RSSM devices on Windows hosts.

Syntax

▶▶—datapath query wwpn—▶▶

Parameters

None

Examples

If you enter the **datapath query wwpn** command, the following output is displayed:

```
[root@abc]> datapath query wwpn
Adapter Name  PortWWN
fscsi0       10000000C925F5B0
fscsi1       10000000C9266FD1
```

datapath remove adapter

The **datapath remove adapter** command dynamically removes an adapter and all paths attached to this adapter from SDD. This command is used when a new adapter must be installed, for example, due to failure of an existing adapter. SDD reserves the last path of a device. If all accessible paths to a device are attached to one adapter, this command fails.

Note: The newly installed adapter can be reclaimed with the **addpaths** command. For detailed procedures, see “addpaths” on page 86.

Note: This command is available for AIX host systems only.

Syntax

```
▶▶—datapath remove adapter—adapter number—▶▶▶▶
```

Parameters

adapter number

The index number of the adapter that you want to remove.

Examples

If you enter the **datapath query adapter** command, the following output is displayed:

```
+-----+
| Active Adapters :4
| Adpt#      Name   State   Mode   Select   Errors  Paths  Active
|  0         fscsi0  NORMAL  ACTIVE 62051    0       10     10
|  1         fscsi1  NORMAL  ACTIVE 65386    3       10     10
|  2         fscsi2  NORMAL  ACTIVE 75697    27      10     10
|  3         fscsi3  NORMAL  ACTIVE 4788     35      10     10
+-----+
```

If you enter the **datapath remove adapter 0** command, the following actions occur:

- The entry for Adpt# 0 disappears from the **datapath query adapter** command output.
- All paths that are attached to adapter 0 disappear from the **datapath query device** command output.

You can enter this command while I/O is running.

```
+-----+
| Active Adapters :3
| Adpt#      Name   State   Mode   Select   Errors  Paths  Active
|  1         fscsi1  NORMAL  ACTIVE 65916    3       10     10
|  2         fscsi2  NORMAL  ACTIVE 76197    27      10     10
|  3         fscsi3  NORMAL  ACTIVE 4997     35      10     10
+-----+
```

The adapter/hard disk Adpt# 0 fscsi0 is removed and the select counts are increased on the other three adapters, indicating that I/O is still running.

datapath remove device path

The **datapath remove device path** command dynamically removes a specific path of an SDD vpath device that is attached to an adapter. SDD reserves the last path of an SDD vpath device. If the requested path is the only accessible path to an SDD vpath device, the command fails.

Note:

1. A removed path can be reclaimed by using the **addpaths** command. See “addpaths” on page 86.
2. This command is available for AIX host systems only.

Syntax

► `datapath remove device—device number—path—path number` ◄

Parameters

device number

The device number shown in the output of the **datapath query device** command.

path number

The path number shown in the output of the **datapath query device** command.

Examples

If you enter the **datapath query device 0** command, the following output is displayed:

```
+-----+
| DEV#:  0  DEVICE NAME: vpath0  TYPE: 2105E20  POLICY: Optimized |
| SERIAL: 20112028 |
|-----+-----+-----+-----+-----+-----+
| Path#  Adapter/Hard Disk  State   Mode    Select  Errors |
|  0     fscsi1/hdisk18      OPEN   NORMAL  557     0 |
|  1     fscsi1/hdisk26      OPEN   NORMAL  568     0 |
|  2     fscsi0/hdisk34      OPEN   NORMAL  566     0 |
|  3     fscsi0/hdisk42      OPEN   NORMAL  545     0 |
+-----+-----+-----+-----+-----+-----+

```

If you enter the **datapath remove device 0 path 1** command, the entry for DEV# 0 Path# 1 (that is, fscsi1/hdisk26) disappears from the **datapath query device 0** command output and the path numbers are rearranged.

```
+-----+
| Success: device 0 path 1 removed |
|-----+-----+-----+-----+-----+-----+
| DEV#:  0  DEVICE NAME: vpath0  TYPE: 2105E20  POLICY: Optimized |
| SERIAL: 20112028 |
|-----+-----+-----+-----+-----+-----+
| Path#  Adapter/Hard Disk  State   Mode    Select  Errors |
|  0     fscsi1/hdisk18      OPEN   NORMAL  567     0 |
|  1     fscsi0/hdisk34      OPEN   NORMAL  596     0 |
|  2     fscsi0/hdisk42      OPEN   NORMAL  589     0 |
+-----+-----+-----+-----+-----+-----+

```

The adapter/hard disk fscsi1/hdisk26 is removed and path 1 is now fscsi0/hdisk34.

The **addpaths** command reclaims the removed path. The mode of the added path is set to NORMAL and its state to either OPEN or CLOSE, depending on the device state.

```
+-----+
| DEV#:  0  DEVICE NAME: vpath0  TYPE: 2105E20  POLICY: Optimized |
| SERIAL: 20112028 |
|-----+-----+-----+-----+-----+-----+
| Path#  Adapter/Hard Disk  State   Mode    Select  Errors |
|  0     fscsi1/hdisk18      OPEN   NORMAL  580     0 |
|  1     fscsi0/hdisk34      OPEN   NORMAL  606     0 |
|  2     fscsi0/hdisk42      OPEN   NORMAL  599     0 |
|  3     fscsi1/hdisk26      OPEN   NORMAL   14     0 |
+-----+-----+-----+-----+-----+-----+

```

Note that fscsi1/hdisk26 is back online with path 3 and is selected for I/O.

datapath set adapter

The **datapath set adapter** command sets all SDD vpath device paths attached to an adapter either to online or offline.

SDD reserves the last path of an SDD vpath device. If all accessible paths to an SDD vpath device are attached to one adapter, the offline option will fail.

Note: The **datapath set adapter offline** command can be entered even when the SDD vpath device is closed.

Syntax

```
▶▶ datapath set adapter adapter number [ online | offline ] ▶▶
```

Parameters

adapter number

The index number of the adapter that you want to change.

online

Sets the adapter online.

offline

Sets the adapter offline.

Examples

If you enter the **datapath set adapter 0 offline** command:

- The mode of Adapter 0 will be changed to OFFLINE while the state of the adapter remains the same.
- All paths attached to adapter 0 change to OFFLINE mode and their states change to Dead, if they were in the Open state.

You can use the **datapath set adapter 0 online** command to cause an adapter that is offline to come online:

- Adapter 0's mode changes to ACTIVE and its state to NORMAL.
- The mode of all paths attached to adapter 0 changes to NORMAL and their state to either OPEN or CLOSE depending on the SDD vpath device state.

datapath set device policy

The **datapath set device policy** command dynamically changes the path selection policy per SDD vpath device.

Syntax

```
▶▶ datapath set device device_num1 [ device_num2 ] policy option ▶▶
```

Parameters

device number1 [*device number2*]

When two device numbers are entered, this command will apply to all the devices whose index numbers fit within the range of these two device index numbers.

option

Specifies one of the following policies:

Windows

- **fo** - failover (no load balance)
- **lb** - least number of I/O and adapter weight
- **sqst** - shortest queue service time
- **lqd** - least queue depth
- **lbytes** - least bytes
- **rr** - round robin
- **rrss** - round robin with subset
- **df** - default (uses lb)

AIX and Linux

- **rr** - round robin
- **rrs** - round robin sequential
- **lb** - load balancing (also known as optimized policy)
- **lbs** - load balancing sequential
- **df** - default policy which is load balancing
- **fo** - failover policy

Note: You can enter the **datapath set device N policy** command to dynamically change the policy associated with SDD vpath devices in either Close or Open state.

Examples

If you enter **datapath set device 2 7 policy rr**, the path-selection policy of SDD vpath devices with device index 2 to 7 is immediately changed to the round robin policy.

datapath set device path

The **datapath set device path** command sets the path of an SDD vpath device either to online or offline. You cannot remove the last path to an SDD vpath device from service. This prevents a data access failure from occurring. The **datapath set device path** command can be entered even when the SDD vpath device is closed.

Syntax

```
▶▶ datapath set device—device number—path—path number——online—  
—offline—▶▶
```

Parameters

device number

The device index number that you want to change.

path number

The path number that you want to change.

online

Sets the path online.

offline

Removes the path from service.

Examples

If you enter the **datapath set device 0 path 0 offline** command, path 0 for device 0 changes to Offline mode.

datapath set performance

The **datapath set performance on** command initiates the performance statistics collection in the SDDDSM driver. The **off** switch stops the statistics collection in the driver.

Note: The switch **on/off** must be used without any device number on the command line.

Note: This command is available for Windows host systems only.

Syntax

```
▶▶ datapath set performance [ -on | -off ] device number <n>
```

Parameters

on/off

Enables or disables performance collection in the SDDDSM driver.

device <n>

The device number for which performance statistics need to be collected.

Examples

If you enter **datapath set performance on**, the following output is generated:

```
Success:Performance Collection set to ON
```

If you enter **datapath set performance device 4**, the following output is generated:

```
Performance Collection set for the required device
```

datapath set qdepth

The **datapath set qdepth** command dynamically sets the queue depth logic of an SDD vpath device to either *enable* or *disable* even when I/O is running.

Note: This command is available for AIX host systems only.

Syntax

► datapath set device _n _m qdepth _enable _disable ►

Parameters

n The beginning vpath number for which the queue depth logic setting is to be applied.

m Then ending vpath number for which the queue depth logic setting is to be applied.

Enable

Enable the queue depth logic.

Disable

Disable the queue depth logic.

Examples

If you enter the **datapath set device 0 2 qdepth disable** command, the following output is displayed:

```
Success: set qdepth_enable to no for vpath0
```

```
Success: set qdepth_enable to no for vpath1
```

```
Success: set qdepth_enable to no for vpath2
```

The `qdepth_enable` ODM attribute of these SDD vpath devices is updated. The following output is displayed when you enter **lsattr -El vpath0**.

```
# lsattr -El vpath0
active_hdisk hdisk66/13AB2ZA1020/fscsi3 Active hdisk False
active_hdisk hdisk2/13AB2ZA1020/fscsi2 Active hdisk False
active_hdisk hdisk34/13AB2ZA1020/fscsi2 Active hdisk False
active_hdisk hdisk98/13AB2ZA1020/fscsi3 Active hdisk False
policy df Scheduling Policy True
pvid 0005f9fdcda4417d0000000000000000 Physical volume identifier False
qdepth_enable no Queue Depth Control True
reserve_policy PR_exclusive Reserve Policy True
serial_number 13AB2ZA1020 LUN serial number False
unique_id yes Device Unique Identification False
```

Appendix A. SDD, SDDPCM, and SDDDSM data collection for problem analysis

The following sections describe enhanced trace capability for SDD, SDDPCM and SDDDSM.

Enhanced trace capability for the SDD and SDDDSM

Starting with SDD 1.5.1.0 (and later), SDD has enhanced its trace capability for problem determination.

SDD has been enhanced to collect SDD trace data periodically and write the trace data to the system local hard drive. SDD maintains 4 files for its trace data:

- `sdd.log`
- `sdd_bak.log`
- `sddsrv.log`
- `sddsrv_bak.log`

Starting with SDDDSM 2.4.3.4 on Windows, the logging feature is enhanced. The size and number of log files that can be maintained on the system are dynamically configurable. The `sddsrv.conf` file has two parameters (**`max_log_size`** and **`max_log_count`**) to configure the size and the number of log files. Trace data is stored in one or all of the following eight files:

- `sdd.log`
- `sdd.log.1`
- `sdd.log.2`
- `sdd.log.3`
- `sddsrv.log`
- `sddsrv.log.1`
- `sddsrv.log.2`
- `sddsrv.log.3`

These files can be found in the following directories:

- AIX - `/var/adm/ras`
- HP-UX - `/var/adm`
- Linux - `/var/log`
- Solaris - `/var/adm`
- Windows Server 2008 and Windows Server 2012 - `%windir%\system32`

Using `sddgetdata` to collect information for problem determination

SDD provides the `sddgetdata` script to collect information used for problem determination. For UNIX platforms, `sddgetdata` creates a tar file or a compressed tar file at the current directory with the current date and time as a part of the file name (for example, `sdddata_hostname_yyyymmdd_hhmmss.tar` or `sdddata_hostname_yyyymmdd_hhmmss.tar.Z`, where `yyymmdd_hhmmss` is the timestamp of the file creation).

For Windows, you can run the **sddgetdata** script from the install directory to collect the data for problem determination. **sddgetdata** creates a cab file in the install directory with the current date and time as part of the file name (for example, `sdddata_hostname_yyyymmdd_hhmmss.cab`), where `yyymmdd_hhmmss` is the timestamp of the file creation). For SDD, the install directory is `%root%\Program Files\IBM\Subsystem Device Driver`. For SDDDSM, the install directory is `%root%\Program Files\IBM\SDDDSM`.

When you report an SDD problem, it is essential to run this script and send this output file for problem determination.

Steps within the **sddgetdata** script might fail depending on the problem and the system condition. In this case, you might have to issue manual commands.

Here is an example output for the AIX platform:

```
/tmp/sdd_getdata>sddgetdata
/tmp/sdd_getdata>ls
./ ../ sdddata_host1_20050315_122521.tar
```

Enhanced trace capability for SDDPCM

For SDDPCM release 2.1.0.7 and later releases, SDDPCM enhanced its trace capability for problem determination. It was enhanced to periodically collect SDDPCM trace data and write the trace data to the system local hard drive. Four files are maintained to store the trace data:

- `pcm.log`
- `pcm_bak.log`
- `pcmsrv.log`
- `pcmsrv_bak.log`

Beginning with SDDPCM 3.0.x.x, SDDPCM will maintain two additional files for its trace data:

- `AE.log`
- `AE_bak.log`

All six of these files can be found in the `/var/adm/ras` directory.

When you report an SDDPCM problem, the following information is essential for problem determination:

- `pcm.log`
- `pcm_bak.log`
- `pcmsrv.log`
- `pcmsrv_bak.log`
- Output of the **pcmpath query adapter** command
- Output of the **pcmpath query device** command
- AIX system error log
- `AE.log`
- `AE_bak.log`

Starting with SDDPCM 2.6.4.x on AIX, the logging feature is enhanced. The size and number of log files that can be maintained on the system are dynamically configurable. The `pcmsrv.conf` file has two parameters (**max_log_size** and

max_log_count) to configure the size and the number of log files. Trace data is stored in one or all of the following eight files:

- pcm.log
- pcm.log.1
- pcm.log.2
- pcm.log.3
- pcmsrv.log
- pcmsrv.log.1
- pcmsrv.log.2
- pcmsrv.log.3

Using **sddpcmgetdata** to collect information for problem determination

SDDPCM provides the **sddpcmgetdata** script to collect information used for problem determination. **sddpcmgetdata** creates a tar file at the current directory with the current date and time as a part of the file name (for example, `sddpcmdata_hostname_yyyymmdd_hhmmss.tar`, where `yyymmdd_hhmmss` is the timestamp of the file creation). When you report an SDDPCM problem, it is essential to run this script and send this tar file for problem determination.

For example:

```
/tmp/sddpcmgetdata>sddpcmgetdata/tmp/sddpcmgetdata>ls  
./ ../ sddpcmdata_test1_20050315_122521.tar
```

Appendix B. System log messages

The following sections list the log messages for AIX SDD and Windows SDDDSM.

AIX error and information messages

This section lists log messages for AIX SDD.

SDD logs error messages into the AIX error log system. To check if SDD has generated an error message, enter the **errpt -a | grep VPATH** command. The **errpt -a | grep SDDSRV** command will show any error messages logged by sddsrv.

The following list shows general SDD error messages.

VPATH_XBUF_NOMEM

An attempt was made to open an SDD vpath device file and to allocate kernel-pinned memory. The system returned a null pointer to the calling program and kernel-pinned memory was not available. The attempt to open the file failed.

VPATH_PATH_OPEN

SDD vpath device file failed to open one of its paths (hdisks). An attempt to open an SDD vpath device is successful if at least one attached path opens. The attempt to open an SDD vpath device fails only when *all* the SDD vpath device paths fail to open.

VPATH_DEVICE_OFFLINE

Several attempts to retry an I/O request for an SDD vpath device on a path have failed. The path state is set to DEAD and the path is taken offline. Enter the **datapath** command to set the offline path to online. For more information, see Chapter 10, "Using the datapath commands," on page 371.

VPATH_DEVICE_ONLINE

SDD supports DEAD path auto_failback and DEAD path reclamation. A DEAD path is selected to send an I/O, after it has been bypassed by 2000 I/O requests on an operational path. If the I/O is successful, the DEAD path is put online, and its state is changed back to OPEN; a DEAD path is put online, and its state changes to OPEN after it has been bypassed by 50 000 I/O requests on an operational path.

VPATH_OUT_SERVICE

An SDD vpath device has no path available for an I/O operation. The state of the SDD vpath device is set to LIMBO. All following I/Os to this SDD vpath device are immediately returned to the caller.

The following list shows the error messages logged by SDD in a persistent reserve environment. See "SDD persistent reserve attributes" on page 56 for more information about persistent reserve.

VPATH_FAIL_RELPRESERVE

An attempt was made to close an SDD vpath device that was not opened with the **RETAIN_RESERVE** option on the persistent reserve. The attempt to close the SDD vpath device was successful; however, the persistent reserve was not released. The user is notified that the persistent reserve is still in effect, and this error log is posted.

VPATH_RESV_CFLICT

An attempt was made to open an SDD vpath device, but the reservation key of the SDD vpath device is different from the reservation key currently in effect. The attempt to open the device fails and this error log is posted. The device could not be opened because it is currently reserved by someone else.

The following are information messages that are logged if you perform AIX Hot Plug procedures with SDD:

VPATH_ADPT_REMOVED

The **datapath remove adapter *n*** command runs. Adapter *n* and its child devices are removed from SDD.

VPATH_PATH_REMOVED

The **datapath remove device *m* path *n*** command runs. Path *n* for device *m* is removed from SDD.

The following error messages are logged by sddsrsv:

SDDSRV_CONF_MISMATCH

This error is logged when sddsrsv finds out hdisk information in the driver is different from what sddsrsv discovered. sddsrsv logs the error to the system error log immediately and every 15 minutes thereafter

SDDSRV_PORTBINDFAIL

This error is logged when sddsrsv cannot bind the TCP/IP port number specified in its sddsrsv.conf file.

SDDSRV_LOG_WFAIL

This error is logged when sddsrsv cannot write its log file (that is, sddsrsv.log) to file system. sddsrsv logs the error to the system error log immediately and every 10 minutes thereafter until sddsrsv can write again.

SDDSRV_DRLOG_WFAIL

This error is logged when sddsrsv cannot write the driver log file (that is, sdd.log) to file system.

SDDSRV_PROBEENABLE

This message is logged when the sddsrsv probing functionality is enabled.

SDDSRV_PROBEDISABLE

This message is logged when the sddsrsv probing functionality is disabled.

SDDSRV_PROBEINTERVAL

This message is logged when the sddsrsv probing interval is changed.

SDDPCM error log messages

This section lists log messages for AIX SDDPCM.

SDDPCM logs error messages into the AIX error log system. To check if SDDPCM generated an error message, enter the **errpt -a | grep SDDPCM** command.

The following list shows general SDD error messages.

SDDPCM_CHKHSACTION

All I/O paths on an active device of an Open HyperSwap device fail, requiring Tivoli Productivity Center for Replication to initiate an unplanned Open HyperSwap.

SDDPCM_OPENPATH_FAILED

One of the SDDPCM MPIO hdisk's paths has failed to open. The failing path is put in the INVALID state if the MPIO hdisk is opened.

SDDPCM_OSPAIRASSOCIATE

Couple a source device and a target device into one Open HyperSwap device.

SDDPCM_OSPAIRBLOCK

An Open HyperSwap device I/O is blocked.

SDDPCM_OSPAIRDISABLED

The Open HyperSwap functionality of an Open HyperSwap device is disabled because there are no available paths on either the source or target devices.

SDDPCM_OSPAIRDISASSOCIATE

Disassociate an Open HyperSwap device from a session. The device is no longer an Open HyperSwap-enabled device

SDDPCM_OSPAIREENABLED

The Open HyperSwap functionality of an Open HyperSwap device is enabled, with paths available to both source and target devices.

SDDPCM_OSPAIRSRFAILED

An Open HyperSwap device failed to perform swap and resume.

SDDPCM_PATH_FAILED

Several attempts to retry an I/O request for an MPIO device on a path have failed, or a path reaches the threshold of continuous I/O errors. The path state is set to FAILED and the path is taken offline. A FAILED path can be automatically recovered by the health checker if the problem is fixed, or the user can enter the **pcmpath set device M path N online** command to manually recover the path. For more information, see "Using SDDPCM pcmpath commands" on page 145.

SDDPCM_PATH_RECOVERED

A failed path is recovered and is in an operational state.

SDDPCM QUIESCETIMEOUT

Exceeded time limit for quiescing I/Os on an Open HyperSwap device.

SDDPCM_RESUMEDONE

Resuming I/Os on all Open HyperSwap devices in a session is complete.

SDDPCM_SESSIONQUIESCE

Quiescing I/Os to all Open HyperSwap devices in a session.

SDDPCM_SESSIONRDY

A session is ready for HyperSwap.

SDDPCM_SESSIONRESUME

Initiate resume I/Os on all Open HyperSwap devices in a session.

SDDPCM_SESSIONSWAPRESUME

Initiate swapping and resuming I/Os to target devices on all Open HyperSwap devices in a session.

SDDPCM_SWAPRESUMEDONE

Swapping and resuming I/Os to target devices on all Open HyperSwap devices in a session is complete.

Windows log messages

This section lists log messages for SDDDSM on a Windows platform.

The log of events can be viewed in Windows Event Viewer under **Windows Logs** > **System**.

The following messages are logged to the System Event log.

Event ID: 1

Event Type: Informational
Event Source: sdddsm
Description: Multipath device X created, identifier I.

Event ID: 2

Event Type: Informational
Event Source: sdddsm
Description: Multipath device X was assigned disk number N by Windows.

Event ID: 3

Event Type: Warning
Event Source: sdddsm
Description: Multipath device X removed.

Event ID: 4

Event Type: Informational
Event Source: sdddsm
Description: Path N added to multipath device X.

Event ID: 5

Event Type: Warning
Event Source: sdddsm
Description: Path N removed from multipath device X by MPIO.

Event ID: 6

Event Type: Informational
Event Source: sdddsm
Description: MPDISK X path N online by driver.

Event ID: 7

Event Type: Informational
Event Source: sdddsm
Description: MPDISK X path N online by sddsrv.

Event ID: 8

Event Type: Informational
Event Source: sdddsm
Description: MPDISK X path N online by datapath.

Event ID: 9

| Event Type: Informational
| Event Source: sdddsm
| Description: MPDISK X path N online by pathtest.

| **Event ID: 10**

| Event Type: Informational
| Event Source: sdddsm
| Description: MPDISK X path N online by unknown application.

| **Event ID: 11**

| Event Type: Warning
| Event Source: sdddsm
| Description: MPDISK X path N offline by driver.

| **Event ID: 12**

| Event Type: Warning
| Event Source: sdddsm
| Description: MPDISK X path N offline by sddsrv.

| **Event ID: 13**

| Event Type: Warning
| Event Source: sdddsm
| Description: MPDISK X path N offline by datapath.

| **Event ID: 14**

| Event Type: Warning
| Event Source: sdddsm
| Description: MPDISK X path N offline by pathtest.

| **Event ID: 15**

| Event Type: Warning
| Event Source: sdddsm
| Description: MPDISK X path N set to dead by driver.

| **Event ID: 16**

| Event Type: Warning
| Event Source: sdddsm
| Description: MPDISK X path N set to dead by sddsrv.

| **Event ID: 17**

| Event Type: Informational
| Event Source: sdddsm
| Description: Adapter X set to online.

| **Event ID: 18**

| Event Type: Warning
| Event Source: sdddsm
| Description: Adapter X set to offline.

| **Event ID: 19**

| Event Type: Informational
| Event Source: sdddsm
| Description: MPDISK X reserved through path N.

| **Event ID: 20**

| Event Type: Warning
| Event Source: sdddsm
| Description: MPDISK X released.

| **Event ID: 21**

| Event Type: Informational
| Event Source: sdddsm
| Description: MPDISK X path Y registered.

| **Event ID: 22**

| Event Type: Warning
| Event Source: sdddsm
| Description: MPDISK X reservation cleared.

| **Event ID: 23**

| Event Type: Warning
| Event Source: sdddsm
| Description: Multipath device X exceeded max paths.

| **Event ID: 100**

| Event Type: Informational
| Event Source: SDD_Service
| Description: The SDD_service was installed.

| **Event ID: 101**

| Event Type: Error
| Event Source: SDD_Service
| Description: The SDD_service cannot be installed.

| **Event ID: 102**

| Event Type: Informational
| Event Source: SDD_Service
| Description: The SDD_service was removed.

| **Event ID: 103**

| Event Type: Error
| Event Source: SDD_Service
| Description: The SDD_service could not be removed.

| **Event ID: 104**

| Event Type: Informational
| Event Source: SDD_Service
| Description: The SDD_service control dispatcher was started.

| **Event ID: 105**

Event Type: Error
Event Source: SDD_Service
Description: The control handler could not be installed.

Event ID: 106

Event Type: Error
Event Source: SDD_Service
Description: The initialization process failed.

Event ID: 107

Event Type: Informational
Event Source: SDD_Service
Description: The SDD_Service service was started.

Event ID: 108

Event Type: Error
Event Source: SDD_Service
Description: The SDD_Service service cannot be started.

Event ID: 109

Event Type: Informational
Event Source: SDD_Service
Description: The SDD_Service service was stopped.

Event ID: 110

Event Type: Informational
Event Source: SDD_Service
Description: The SDD_Service service cannot be stopped.

Event ID: 111

Event Type: Error
Event Source: SDD_Service
Description: The SDD_Service service cannot be bind to port.

Event ID: 112

Event Type: Error
Event Source: SDD_Service
Description: The SDD_Service service encountered a configuration mismatch.

Event ID: 113

Event Type: Error
Event Source: SDD_Service
Description: The SDD_Service failed to write the server log to file.

Event ID: 114

Event Type: Error
Event Source: SDD_Service
Description: The SDD_Service failed to write the driver log to file.

Event ID: 115

| Event Type: Informational
| Event Source: SDD_Service
| Description: The SDD_Service service probing enabled with interval T min.

| **Event ID: 116**

| Event Type: Informational
| Event Source: SDD_Service
| Description: The SDD_Service service probing disabled.

| **Event ID: 117**

| Event Type: Informational
| Event Source: SDD_Service
| Description: The SDD_Service service probing interval changed to T min.

Appendix C. SDDDSM tools

Several support utilities are available to help you diagnose problems or issues that you might face when you work with SDDDSM.

Overview

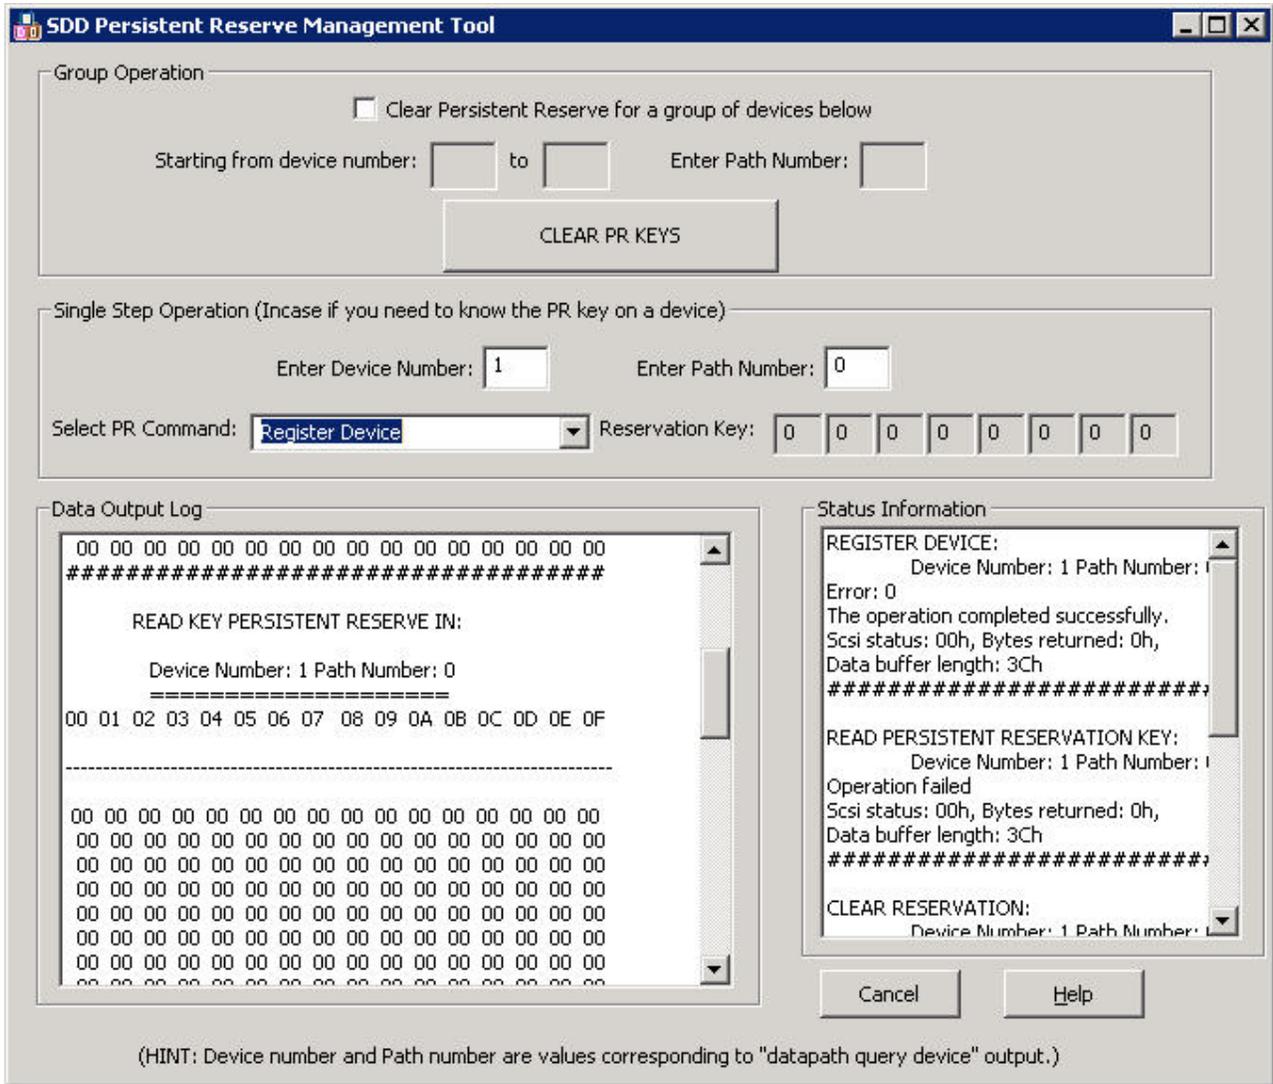
The following support utilities or tools are packaged with the SDDDSM release:

- PRTool.exe
- sddutils.exe
- HBAInfo.exe
- gethba.exe
- sddgetdata.bat
- sasflash.exe
- saswwpn.bat
- pathtest.exe
- sddsrv.exe
- mpiolist.exe

All these tools are available in the installation directory of SDDDSM.

PRTool.exe

Use this tool to read or set data that is related to persistent reservation on the disk. You can read registrations, read reservations, or set registrations and reservations on the disk.



Choose one of the following methods:

- Group operation - Use this method if the persistent reservation is to be removed from a bunch of devices.
- Single-step operation - Use this method for clearing persistent reservation from a device, one at a time.

Only one method can be used at a time. The tool is tightly integrated with SDD "datapath" utility. Here the "Device Number" and "Device Path" are used with "DEV#" and "Path#" from the "datapath query device" command. For more information, see the PRTool.chm file, which is available in the SDDDSM installation directory.

sddutils.exe

Use this tool to save the system event log in the local directory. The log is saved in .evt format.

```
C:\Program Files\IBM\SDDDSM>sddutils.exe getsystemlog
Saved the system event log

C:\Program Files\IBM\SDDDSM>dir *.evt
Volume in drive C has no label.
Volume Serial Number is E442-59FF

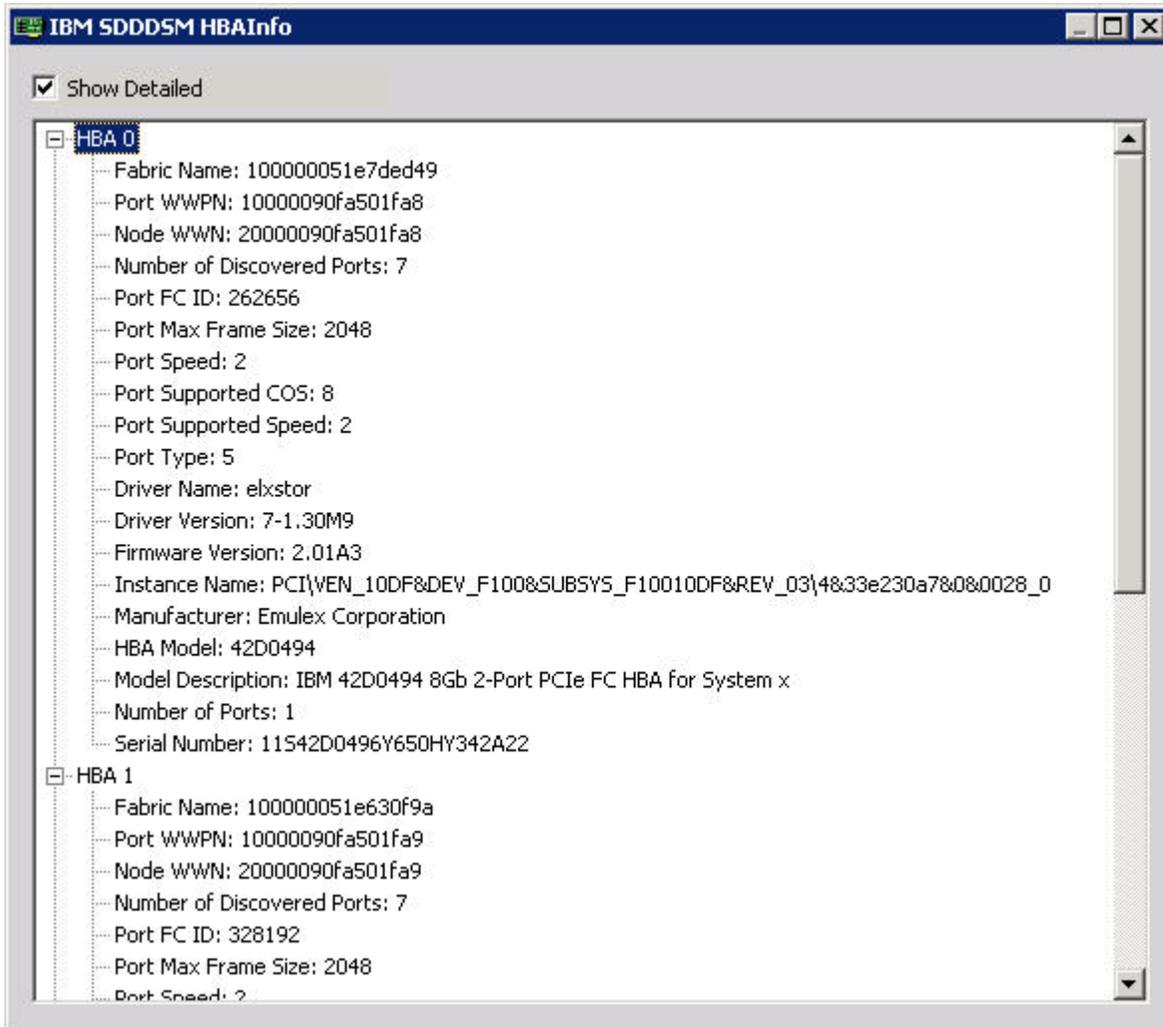
Directory of C:\Program Files\IBM\SDDDSM

11/13/2014  08:12 PM                20,975,616 System.evt
             1 File(s)                20,975,616 bytes
             0 Dir(s)  213,670,576,128 bytes free

C:\Program Files\IBM\SDDDSM>_
```

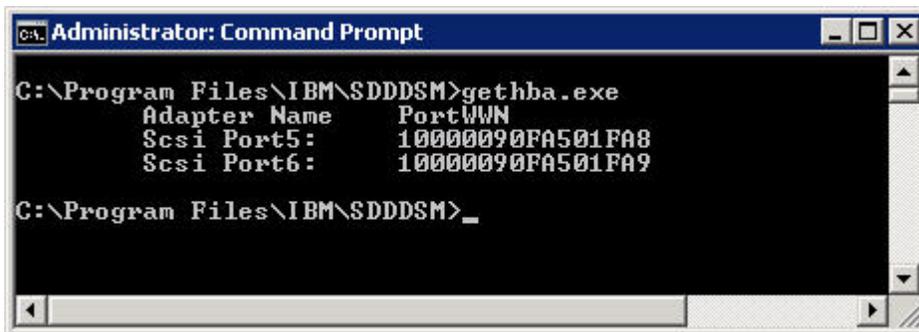
HBAInfo.exe

Use this tool to get detailed HBA information. You must install the VC Redistributable files on host.



gethba.exe

Use this tool to get Qlogic/Emulex HBA WWPNs. Alternatively, you can use datapath query wpn or use the HBAInfo.exe file.

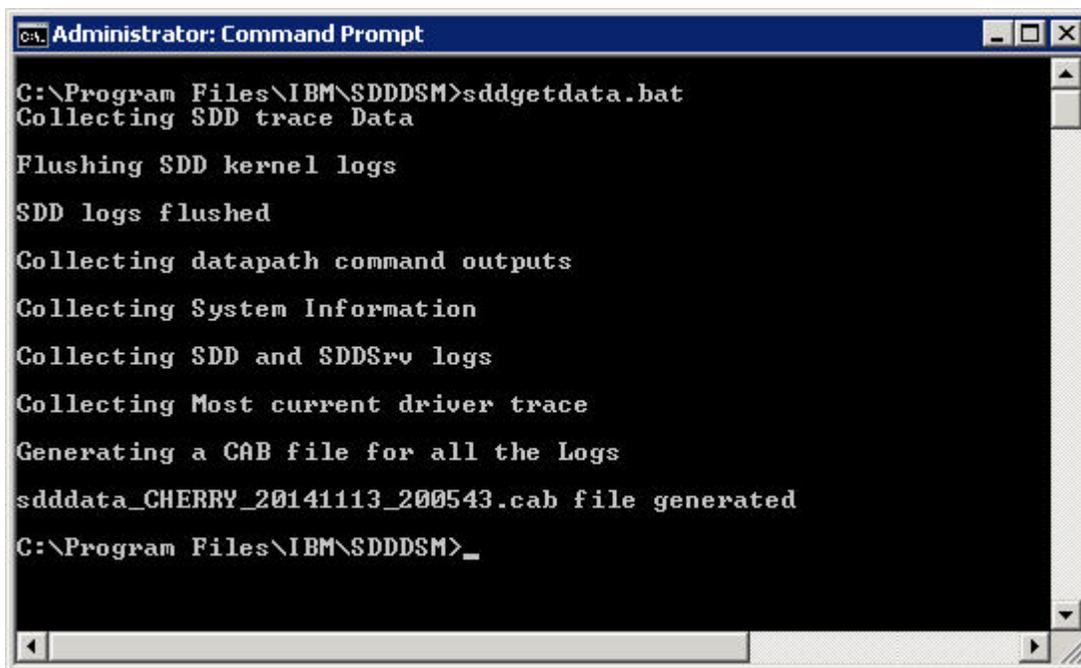


sddgetdata.bat

Use this tool to get SDD trace data.

This tool creates a cab file in the installation directory with the current date and time as part of the file name. For example, sdddata_hostname_yyyymmdd_hhmmss.cab, where yyyymmdd_hhmmss is the time stamp of the file creation. The cab file includes the following information:

- systemInfo
- "HKLM \SYSTEM\CurrentControlSet", "HKLM\HARDWARE\DEVICEMAP", and "HKLM\Cluster" output from registry
- sdddsm directory contents
- hba details
- datapath outputs
- pathtest trace
- sddsrv logs
- cluster logs



```
Administrator: Command Prompt
C:\Program Files\IBM\SDDDSM>sddgetdata.bat
Collecting SDD trace Data
Flushing SDD kernel logs
SDD logs flushed
Collecting datapath command outputs
Collecting System Information
Collecting SDD and SDDSRV logs
Collecting Most current driver trace
Generating a CAB file for all the Logs
sdddata_CHERRY_20141113_200543.cab file generated
C:\Program Files\IBM\SDDDSM>_
```

sasflash.exe

Use this tool to get SAS HBA details and to download the firmware.

You can use the following parameters when you run the sasflash.exe file:

- <-c x > - x stands for adapter index
- <-verbose> - set output to verbose mode
- <-l x> - log output to a file
- <-f x> - flash firmware image
- "x" - firmware file to flash

Important: For the current version of the sasflash.exe file and for information about the current command-line options, go to the LSI website (<http://technet.microsoft.com/en-us/library/dd361921.aspx>).

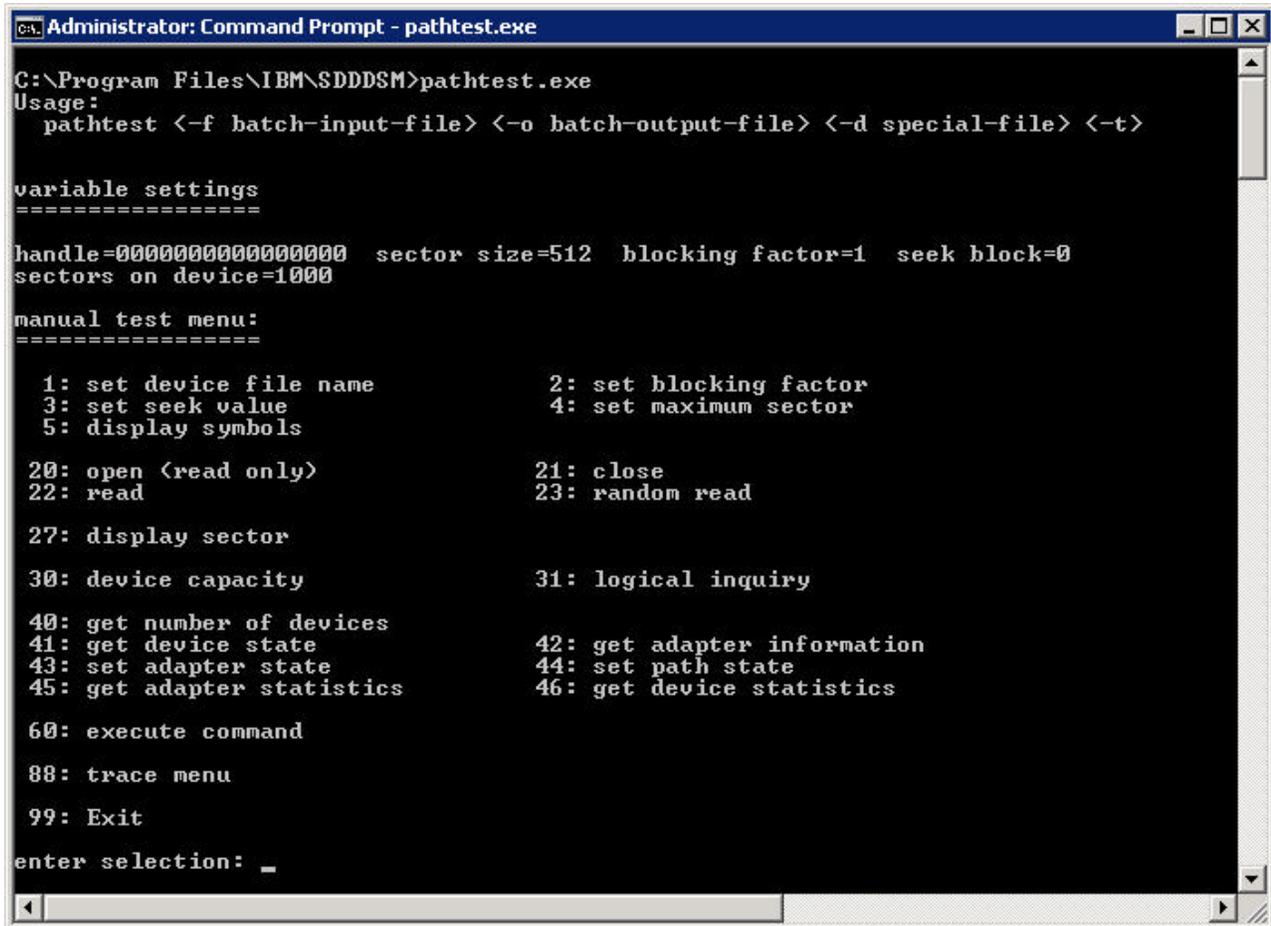
saswwpn.bat

Use this batch file to gather information about SAS HBA. This tool uses the sasflash.exe tool.

```
sasflash.exe -o -listsasadd -c %CONTROLLER% >NUL
```

pathtest.exe

Use this tool to start tracing manually. The following figure shows a sample output.



```
Administrator: Command Prompt - pathtest.exe
C:\Program Files\IBM\SDDDSM>pathtest.exe
Usage:
  pathtest <-f batch-input-file> <-o batch-output-file> <-d special-file> <-t>

variable settings
=====

handle=0000000000000000  sector size=512  blocking factor=1  seek block=0
sectors on device=1000

manual test menu:
=====

  1: set device file name          2: set blocking factor
  3: set seek value                4: set maximum sector
  5: display symbols

20: open (read only)              21: close
22: read                          23: random read

27: display sector

30: device capacity               31: logical inquiry

40: get number of devices         42: get adapter information
41: get device state              44: set path state
43: set adapter state             46: get device statistics
45: get adapter statistics

60: execute command

88: trace menu

99: Exit

enter selection: _
```

sddsrv.exe

sddsrv.exe is an SDDDSM path recovery daemon that runs as a service in the Windows OS.

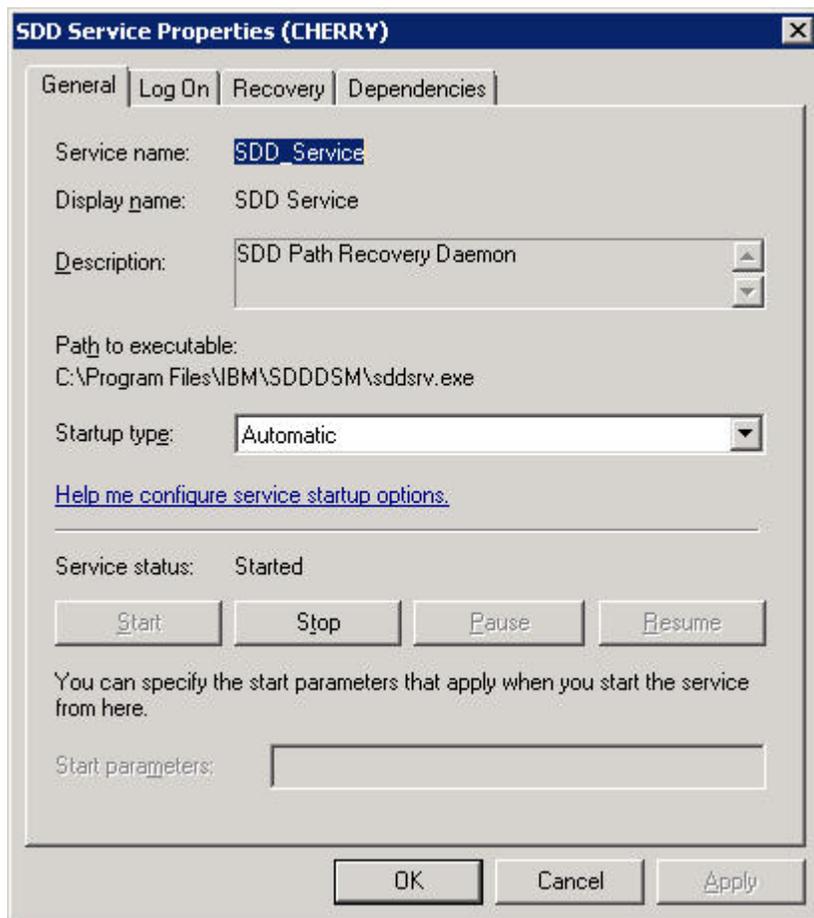
This tool generates log files, sdd.log and sddsrv.log, in the windows\system32 folder to help with debugging any issues. The log files get generated even for the release builds of the sdddsm driver. The sddgetdata.bat file copies these logs in to its cab file.

To use this tool, run the following command:

```
sddsrv d|v|q|f|s0|s1 \n
```

where

- d = produce detailed console output
- v = produce summary console output
- q = produce no console output (default)
- f = allow multiple servers on the same host
- s0 = disable driver interface
- s1 = enable driver interface



The sample content of the sdd.log file is as follows:

```
GmSrbDeviceControl(1126) MPD-000-60050768019B02A8100000000000192, SCSI_ADDRESS 05-00-00-00
GmSrbDeviceControl(1126) MPD-001-60050768019B02A8100000000000193, SCSI_ADDRESS 05-00-00-01
GmSrbDeviceControl(1126) MPD-002-60050768019B02A8100000000000180, SCSI_ADDRESS 05-00-00-02
GmClusterPRIN(1579) MPD-002-60050768019B02A8100000000000180, Key Count= 8, Ac-tivePathCount 8
GmClusterPRIN(1579) MPD-000-60050768019B02A8100000000000192, Key Count= 8, Ac-tivePathCount 8
```

The sample content of the sddsrsv.log file is as follows:

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*

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan*

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 publications. IBM may make improvements and/or changes in the product(s) and/or 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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
Information Enabling Requests
Dept. DZWA
5600 Cottle Road
San Jose, CA 95193
U.S.A.*

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

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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 measurements 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.

COPYRIGHT LICENSE: This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol ([®] or [™]), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Java[™] is a trademark of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

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

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

IBM license agreement for machine code

Read Before Using

IMPORTANT

YOU ACCEPT THE TERMS OF THIS IBM LICENSE AGREEMENT FOR MACHINE CODE BY YOUR USE OF THE HARDWARE PRODUCT OR MACHINE CODE. PLEASE READ THE AGREEMENT CONTAINED IN THIS BOOK BEFORE USING THE HARDWARE PRODUCT.

Licensee accepts the terms of this IBM License Agreement for Machine Code ("License Agreement") by Licensee's initial use of a Machine that contains Machine Code (including Licensed Internal Code), whether the rightful possessor of a Machine (the "Licensee") obtains such Machine from IBM or from a third party.

This License Agreement applies to all Machine Code provided to Licensee for such Machine and replaces any license agreement previously provided to Licensee for such Machine Code. This License Agreement also replaces the IBM Agreement for Licensed Internal Code. For Licensed Internal Code Licensee acquires after September 30, 2008, any reference to the "IBM Agreement for Licensed Internal Code" in an agreement between Licensee and IBM means this License Agreement.

1. Definitions

Built-in Capacity -- any computing resource or capability that may be included with a Machine and that is to remain inactive, or for which use is restricted, until the right to access and use the resource or capability is properly acquired directly from IBM or through an authorized IBM reseller. Such computing resources and capabilities include, without limitation, processors, memory, storage, interactive processing capacity, and workload-specific resources or capabilities (such as a specific operating system, programming language, or application to which use of the Machine is limited).

Licensed Internal Code (also referred to as "LIC") -- machine Code used by certain Machines that IBM or an IBM reseller identifies to Licensee as a "Specific Machine".

Machine -- a hardware device, its features, conversions, upgrades, elements or accessories, or any combination of them.

Machine Code -- microcode, basic input/output system code ("BIOS"), utility programs, device drivers, diagnostics, and any other code (all subject to the exclusions in this License Agreement) delivered with an Machine for the purpose of enabling the Machine's function(s) as stated in the IBM document entitled "Official Published Specifications" applicable to such Machine. Machine Code does not include programs and code provided under open source licenses or other separate license agreements. The term "Machine Code" includes LIC, any whole or partial copy of Machine Code, and any fix, patch, or replacement provided for Machine Code (including LIC).

2. License

International Business Machines Corporation, one of its subsidiaries, or a third party owns (including, without limitation, ownership of all copyrights in) Machine Code and all copies of Machine Code (including, without limitation, copies of the original Machine Code and copies made from copies). Machine Code is copyrighted and licensed (not sold). IBM licenses Machine Code to only one rightful possessor at a time.

2.1 Authorized Use

IBM grants the Licensee a nonexclusive license to use Machine Code on, or in conjunction with, only the Machine for which IBM provided it, and only to the extent of IBM authorizations that Licensee has acquired for access to and use of Built-in-Capacity. If Licensee's use of Built-in-Capacity exceeds such IBM authorizations, Licensee agrees to pay IBM or (if applicable) an authorized IBM reseller the full price of permanent, unrestricted use of the Built-in-Capacity at the then current price. Licensee is not authorized to use such Built-in-Capacity until such payment is made.

Under each license, IBM authorizes Licensee to do only the following:

- a. execute Machine Code to enable the Machine to function according to its Official Published Specifications;
- b. use only the Built-in-Capacity that Licensee has properly acquired for the Machine directly from IBM or through an authorized IBM reseller;

- c. make a reasonable number of copies of Machine Code to be used solely for backup or archival purposes, provided i) Licensee reproduces the copyright notice and any other legend of ownership on any such copies and ii) uses the copies only to replace the original, when necessary; and
- d. execute and display Machine Code as necessary to maintain the Machine.

No other licenses or rights (including licenses or rights under patents) are granted either directly, by implication, or otherwise.

2.2 Actions Licensee May Not Take

Licensee agrees to use Machine Code only as authorized above. Licensee may not do any of the following:

- a. otherwise copy, display, transfer, adapt, modify, or distribute (electronically or otherwise) Machine Code, except as IBM may authorize in a Machine's user documentation or in writing to Licensee;
- b. reverse assemble, reverse compile, otherwise translate, or reverse engineer Machine Code unless expressly permitted by applicable law without the possibility of contractual waiver;
- c. assign the license for Machine Code; or
- d. sublicense, rent, or lease Machine Code or any copy of it.

2.3 Replacements, Fixes, and Patches

Licensee agrees to acquire any replacement, fix or patch for, or additional copy of, Machine Code directly from IBM in accordance with IBM's standard policies and practices. Unless Licensee is provided with a different IBM Machine Code license agreement, the terms of this License Agreement will apply to any replacement, fix or patch for, or additional copy of, Machine Code Licensee acquires for the Machine. If such additional or different terms apply to any replacement, fix, or patch, Licensee accepts them when Licensee downloads or uses the replacement, fix, or patch.

2.4 Machine Code Transfers

Licensee may transfer possession of Machine Code and its media to another party only with the transfer of the Machine for which that Machine Code is authorized. In the event of such transfer, Licensee agrees to 1) destroy all of Licensee's copies of that Machine Code that were not provided by IBM, 2) either provide to the other party all Licensee's IBM-provided copies of Machine Code or destroy them, 3) provide to the other party a copy of this License Agreement, and 4) provide to the other party all user documentation. IBM licenses the other party to use Machine Code when that party accepts the terms of this License Agreement and is the rightful possessor of the associated Machine.

2.5 Termination

Licensee's license for Machine Code terminates when Licensee no longer rightfully possesses the associated Machine.

3. Built-in Capacity

Built-in-Capacity is limited by certain technological measures in Machine Code. Licensee agrees to IBM's implementation of such technological measures to protect Built-in-Capacity, including measures that may impact availability of data or performance of the Machine. As a condition of Licensee's license to use Machine Code under this License Agreement, Licensee may not (i) circumvent such technological measures or use a third party or third party product to do so, or (ii) otherwise access or use unauthorized Built-in-Capacity. If IBM determines that changes are necessary to the technological measures designed to limit access to, or use of, Built-in-Capacity, IBM may provide Licensee with changes to such technological measures. As a condition of Licensee's license to use Machine Code under this License Agreement, Licensee agrees, at IBM's option, to apply or allow IBM to apply such changes.

4. Relationship to Other Agreements

If Licensee obtained Machine Code directly from IBM and an IBM Customer Agreement ("ICA") or an equivalent agreement is in effect between Licensee and IBM, the terms of this License Agreement are incorporated by reference into the ICA or the applicable equivalent agreement.

If no ICA or equivalent agreement is in effect between Licensee and IBM or Licensee obtained Machine Code through an IBM reseller or other third party, the IBM Statement of Limited Warranty ("SOLW") is incorporated by reference into this License Agreement and applies to Machine Code.

To the extent of any conflict between the terms of this License Agreement and those of (i) the ICA or applicable equivalent agreement or (ii) the SOLW, the terms of this License Agreement prevail.

Glossary

This glossary includes terms for disk storage system products and Virtualization 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: www-01.ibm.com/software/globalization/terminology/index.jsp. Definitions derived from this book 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 being developed by ISO/IEC JTC1/SC1 have the symbol (T) after the definition, indicating that final agreement has not been reached among the participating National Bodies of SCI.

This glossary uses the following cross-reference forms:

See This 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

This refers the reader to one or more related terms.

Special characters

1750 The machine type for the IBM System Storage DS6000 series. Models for the DS6000 include the 511 and EX1.

1820 The machine number for the RSSM.

2105 The machine number for the IBM TotalStorage Enterprise Storage Server (ESS). Models of the ESS are expressed as the number 2105 followed by “Model <xxx>”, such as 2105 Model 800. The 2105 Model 100 is an ESS expansion enclosure that is typically referred to simply as the Model 100. See also *IBM TotalStorage Enterprise Storage Server* and *Model 100*.

2107 A hardware machine type for the IBM System Storage DS8000 series. Hardware models for the 2107 include base units 921, 922, 931, 932, 9A2, 9B2 and expansion units 92E and 9AE.

2145 A hardware machine type for the IBM System Storage SAN Volume Controller. Models of the SAN Volume Controller are expressed as the number 2145 followed by “-xxx”, such as 2145-8G4. Hardware models for the 2145 include 2145-4F2, 2145-8F2, 2145-8F4, and 2145-8G4.

3390 The machine number of an IBM disk storage system. The ESS, when interfaced to IBM S/390 or IBM System z® 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.

8-pack See *disk eight pack*.

/ file system

The root file system; contains files that contain machine-specific configuration data.

/tmp file system

A shared storage location for files.

/usr file system

Contains files and programs necessary for operating the machine.

/var file system

Contains files that are variable on a per-client basis, such as spool and mail files.

A

access To obtain the use of a computer resource.

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 disk storage system product during initial configuration. It enables all fibre-channel-attached host systems with no defined access profile to access all logical volumes on the disk storage system. With a profile defined in ESS Specialist 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 acknowledgement and acknowledgement*.

active Copy Services server

The Copy Services server that manages the Copy Services domain. Either the primary or the backup Copy Services server can be the active Copy Services server. The backup Copy Services server is available to become the active Copy Services server if the primary Copy Services server fails. See also *backup Copy Services server*, *Copy Services client*, and *primary Copy Services server*.

active/active mode

A configuration that enables one controller node of a storage system pair to process I/O requests and provide a standby capability for the other controller node. Generally, an active/active storage system involves a battery-backed mirrored cache, in which the cache content of a controller is mirrored to another for data integrity and availability.

active/passive mode

A configuration that enables one controller node of a storage system pair to process I/O requests, while the other controller node is idle in standby mode

ready to take over I/O activity if the active primary controller fails or is taken offline.

alert A message or log that a storage facility generates as the result of error event collection and analysis. An alert indicates that a service action is required.

allegiance

In Enterprise Systems Architecture/390, 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

In a disk storage system, the space that is allocated to volumes but not yet assigned. See also *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 ESS Specialist, the label on an icon that represents all connections that are using fibre-channel adapters between the ESS and hosts and that are not completely defined to the ESS. See also *anonymous host*, *pseudo-host*, and *access-any mode*.

anonymous host

Synonym for "pseudo-host" (in contrast to the Anonymous label that appears on some pseudo-host icons. See also *Anonymous* and *pseudo-host*.

ANSI See *American National Standards Institute*.

APAR See *authorized program analysis report*. (GC)

arbitrated loop

For fibre-channel connections, a 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 are used to define logical volumes or devices. More specifically, regarding the disk storage system, an array is a group of disks designated by the user to be

managed by the RAID-5 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

On a disk storage system, the space allocated to a volume and assigned to a port.

authorized program analysis report (APAR)

A report of a problem caused by a suspected defect in a current, unaltered release of a program. (GC)

availability

The degree to which a system or resource is capable of performing its normal function. See *data availability*.

B

backup Copy Services server

One of two Copy Services servers in a Copy Services domain. The other Copy Services server is the primary Copy Services server. The backup Copy Services server is available to become the active Copy Services server if the primary Copy Services server fails. A Copy Services server is software that runs in one of the two clusters of an ESS, and manages data-copy operations for that Copy Services server group. See also *active Copy Services server*, *Copy Services client*, and *primary Copy Services server*.

bay

In the disk storage system, the physical space used for installing SCSI, ESCON, and fibre-channel host adapter cards. The ESS has four bays, two in each cluster. See also *service boundary*.

bit

Either of the digits 0 or 1 when used in the binary numeration system. (T)

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)

In the disk storage system, 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 A group of eight adjacent binary digits that represent one EBCDIC character.

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

In the disk storage system, 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. The opposite of *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 facility.

cache miss

An event that occurs when a read operation is sent to the cluster, but the data is not found in cache. The opposite of *cache hit*.

call home

A communication link established between the disk storage system and a service provider. The disk storage system 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

Connecting network controllers to each other in a succession of levels, to concentrate many more lines than a single level permits.

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 a disk storage system sends to it.

CCR See *channel command retry*.

CCW See *channel command word*.

CD See *compact disc*.

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 disk storage system code and documentation.

CEC See *computer-electronic complex*.

channel

In Enterprise Systems Architecture/390, 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)

In Enterprise Systems Architecture/390, 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)

In Enterprise Systems Architecture/390, a

data structure that specifies an I/O operation to the channel subsystem.

channel path

In Enterprise Systems Architecture/390, the interconnection between a channel and its associated control units.

channel subsystem

In Enterprise Systems Architecture/390, the part of a host computer that manages I/O communication between the program and any attached control units.

channel-subsystem image

In Enterprise Systems Architecture/390, 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

In the disk storage system, a partition capable of performing all disk storage system functions. With two clusters in the disk storage system, any operational cluster can take over the processing of a failing cluster.

In the AIX operating system, a group of nodes within a complex.

cluster processor complex (CPC)

In the disk storage system, the unit within a cluster that provides the management function for the disk storage system. It consists of cluster processors, cluster memory, and related logic.

Code Distribution and Activation (CDA)

Process of installing licensed machine code on a disk storage system while applications continue to run.

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, UNIX shell commands). See also *Copy Services command-line interface*.

compression

The process of eliminating gaps, empty fields, redundancies, and unnecessary data to shorten the length of records or blocks.

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 download of licensed machine code

Process of installing licensed machine code 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 facility provides for this function.

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, such as can be provided by a personal computer. See also *IBM TotalStorage ESS Master Console*.

contingent allegiance

In Enterprise Systems Architecture/390, 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 unit (CU)

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.

In Enterprise Systems Architecture/390, a storage server with ESCON, FICON[®], or OEMI interfaces. The control unit adapts a native device interface to an I/O interface supported by an ESA/390 host system.

In the ESS, the portion of the ESS that supports the attachment of emulated CKD devices over ESCON, FICON, or OEMI interfaces. See also *cluster*.

control-unit image

In Enterprise Systems Architecture/390, a logical subsystem that is accessed through an ESCON or FICON 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 ESS can emulate 3990-3, TPF, 3990-6, or 2105 control units.

control-unit initiated reconfiguration (CUIR)

A software mechanism that the ESS uses to request that an operating system of an IBM System z or S/390 host verify that one or more subsystem resources can be taken offline for service. The ESS 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 client

Software that runs on each ESS 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.

See also *active Copy Services server*, *backup Copy Services server*, and *primary Copy Services server*.

Copy Services CLI

See *Copy Services Command-Line Interface*.

Copy Services domain

A collection of user-designated ESS clusters participating in Copy Services functions managed by a designated active Copy Services server. See also *Copy Services server*, *dual-active server*, and *single-active server*.

Copy Services command-line interface (Copy Services CLI)

In the ESS, command-line interface software provided with ESS Copy Services and used for invoking Copy Services functions from host systems attached to the ESS. See also *command-line interface*.

Copy Services server

An ESS cluster designated by the copy services administrator to perform the ESS Copy Services functions. See also *active Copy Services server*, *backup Copy Services server*, and *primary Copy Services server*.

Copy Services server group

A collection of user-designated ESS clusters participating in Copy Services functions managed by a designated active Copy Services server. A Copy Services server group is also called a Copy Services domain. See also *active Copy Services server*, *backup Copy Services server*, and *primary Copy Services server*.

copy set

A set of volumes that contain copies of the same data. All the volumes in a copy set are the same format (count key data [CKD] or fixed block) and size.

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)

In Enterprise Systems Architecture/390, a data-record format employing self-defining record formats in which each record is represented by up to three fields: a *count* field identifying the record and specifying its format, an optional *key* field that can be used to identify the data area contents, and an optional *data* field that typically contains the user data. For CKD records on the ESS, 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

In the ESS, a volume in count-key-data (CKD) format that is not a standard volume, which basically 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 S/390 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 *SSA 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

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 S/390 and IBM System z storage on an ESS, 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

normally used by the operating system to manage the track. See also *count-key-data* and *fixed-block architecture*.

data sharing

The ability of multiple host systems to concurrently utilize data that they store on one or more storage devices. The storage facility 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 eight pack*.

dedicated storage

Storage within a storage facility 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 ESS must destage logical data units with active write units before they can be demoted.

destaging

Movement of data from an online or higher priority to an offline or lower priority device. The ESS stages incoming data into cache and then destages it to disk.

device In Enterprise Systems Architecture/390, a disk drive.

device adapter (DA)

A physical component of the ESS that provides communication between the clusters and the storage devices. The ESS has eight device adapters that it deploys in pairs, one from each cluster. DA pairing enables the ESS to access any disk drive from either of two paths, providing fault tolerance and enhanced availability.

device address

In Enterprise Systems Architecture/390, the field of an ESCON or FICON device-level frame that selects a specific device on a control-unit image.

device ID

In the ESS, 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 DDMs.

device number

In Enterprise Systems Architecture/390, 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 DDM to a spare DDM. The subsystem maintains data access during the process.

DFS See *distributed file service*.

direct access storage device (DASD)

A mass storage medium on which a computer stores data.

A disk device.

disk cage

A container for disk drives. Each disk cage supports eight disk eight packs (64 disks).

disk drive

Standard term for a disk-based nonvolatile storage medium. The ESS 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 eight pack*.

disk eight pack

In the ESS, a group of eight disk drive modules (DDMs) installed as a unit in a DDM bay.

disk group

In the ESS, a collection of disk drives in the same SSA loop set up by the ESS to

be available to be assigned as a RAID array. A disk group can be formatted as CKD or fixed block, and as RAID or non-RAID, or it can be left unformatted. A disk group is a logical assemblage of eight disk drives, in contrast to a disk eight pack. See also *disk eight pack*.

disk storage system

One or more storage devices that are installed with a storage software application to provide a single common pool of storage that is used to store, safeguard, retrieve, and share data. Most disk storage systems also include disaster planning and recovery options. In SDD, a disk storage system refers to an ESS, DS6000, or DS8000 device.

distributed file service (DFS)

A service that provides data access over IP networks.

DNS See *domain name system*.

domain

That part of a computer network in which the data processing resources are under common control.

In TCP/IP, the naming system used in hierarchical networks.

A Copy Services server group, in other words, the set of clusters designated by the user 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.

drawer

A unit that contains multiple DDMs and provides power, cooling, and related interconnection logic to make the DDMs accessible to attached host systems.

drive

A peripheral device, especially one that has addressed storage media. See also *disk drive module*.

The mechanism used to seek, read, and write information on a storage medium.

dual-active mode

With respect to a Copy Services server,

- the mode of operation of the server when the LIC level of the associated ESS cluster is 2.0 or higher.
- With respect to a Copy Services domain, the mode of operation of the domain, when the Copy Services servers are dual-active servers. See also *Copy Services server*, *Copy Services domain*, *mixed mode*, and *single-active server*.
- duplex**
Regarding ESS Copy Services, the state of a volume pair after PPRC has completed the copy operation and the volume pair is synchronized.
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 ESS. See also *F10*.
- E20** The predecessor of the F20 model of the ESS. 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 eight 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 CKD 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 Storage Server**
See *IBM TotalStorage Enterprise Storage Server*.
- 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)**
An Enterprise Systems Architecture/390 and IBM System z computer peripheral interface. The I/O interface uses ESA/390 logical protocols over a serial interface that configures attached units to a communication fabric.
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*.
- error-recording data set (ERDS)**
On S/390 and IBM System z hosts, a data set that records data-storage and data-retrieval errors. A service information message (SIM) provides the error information for the ERDS.
- ERP** See *error recovery procedure*.
- 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

An S/390 or IBM System z channel that supports ESCON protocols.

ESCON director (ESCD)

An I/O interface switch that provides for the interconnection of multiple ESCON interfaces in a distributed-star topology.

ESCON host systems

S/390 or IBM System z hosts that attach to the ESS 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)

In Enterprise Systems Architecture/390, a function that enables LPARs to share an ESCON channel path by providing each LPAR with its own channel-subsystem image.

EsconNet

In ESS Specialist, 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 ESS. See also *pseudo-host* and *access-any mode*.

ESD See *electrostatic discharge*.

eServer

See *IBM eServer*.

ESS See *IBM TotalStorage Enterprise Storage Server*.

ESS Copy Services

In the ESS, a collection of optional software features, with a Web-browser interface, used for configuring, managing, and monitoring data-copy functions.

ESS Copy Services CLI

See *Copy Services Command-Line Interface*.

ESS Expert

See *IBM TotalStorage Enterprise Storage Server Expert*.

ESS Master Console

See *IBM TotalStorage ESS Master Console*.

ESSNet

See *IBM TotalStorage Enterprise Storage Server Network*.

ESS Specialist

See *IBM TotalStorage Enterprise Storage Server Specialist*.

Expert See *IBM TotalStorage Enterprise Storage Server Expert*.

extended binary-coded decimal interchange code (EBCDIC)

A coding scheme developed by IBM 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 CKD architecture.

Extended Remote Copy (XRC)

A function of a storage server that assists a control program to maintain a consistent copy of a logical volume on another storage facility. 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.

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*.

F

F10 A model of the ESS featuring a single-phase power supply. It has fewer expansion capabilities than the Model F20.

F20 A model of the ESS featuring a three-phase power supply. It has more expansion capabilities than the Model F10, including the ability to support a separate expansion enclosure.

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

Cluster recovery from failover following repair. See also *failover*.

failover

In SAN Volume Controller, the function that occurs when one redundant part of the system takes over the workload of another part of the system that has failed.

In the ESS, the process of transferring all control of the ESS to a single cluster in the ESS when the other cluster in the ESS fails. See also *cluster*.

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 ESS 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 ESS 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. See American National Standards Institute (ANSI) X3T11/93-275. In this topology, two or more fibre-channel end points are interconnected through a looped interface. The ESS supports this topology.

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 ESA/390 command protocols). See American National Standards Institute (ANSI) X3.230-199x. See also *fibre-channel protocol*.

FICON

See *fibre-channel connection*.

FiconNet

In ESS Specialist, 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 ESS. 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 for the ESS that can make an instant copy of data, that is, a point-in-time copy of a volume.

FRU See *field replaceable unit*.

FTP See *File Transfer Protocol*.

full duplex

See *duplex*.

G

GB See *gigabyte*.

GDPS®

See *Geographically Dispersed Parallel Sysplex®*.

Geographically Dispersed Parallel Sysplex™ (GDPS)

An S/390 multisite application-availability solution.

gigabyte (GB)

A gigabyte of storage is 10⁹ bytes. A gigabyte of memory is 2³⁰ bytes.

group In ESS documentation, a nickname for two different kinds of groups, depending on the context. See *disk eight pack* or *Copy Services server group*.

H

HA See *host adapter*.

HACMP

See *High-Availability Cluster Multiprocessing*.

hard disk drive (HDD)

A storage medium within a storage server used to maintain information that the storage server requires.

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 IBM System i® 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.

HBA See *host bus adapter*.

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 ESS. 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

A function provided by storage management software such as Tivoli Storage Management or Data Facility Storage Management Subsystem/MVS (DFSMS/MVS) to automatically manage free space based on the policy that the storage administrator sets.

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 Multiprocessing (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 (HA)

A nine-byte field at the beginning of a track that contains information that identifies the physical track and its association with a cylinder.

Note: In the ESS, the acronym HA is shared between home address and host adapter. See also *host adapter*.

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 (HA)

A physical subunit of a storage server that provides the ability to attach to one or more host I/O interfaces. The Enterprise Storage Server has four HA bays, two in each cluster. Each bay supports up to four host adapters.

In the ESS, the acronym HA is shared between home address and host adapter. See also *home address*.

host bus adapter

An interface card that connects a host bus, such as a peripheral component interconnect (PCI) bus, to the storage area network.

host name

The Internet address of a machine in the network. In the ESS, 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 (S/390 or IBM system z) or of the open-systems type, that is connected to the ESS. S/390 or IBM System z hosts are connected to the ESS through ESCON or FICON interfaces. Open-systems hosts are connected to the ESS 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

IBM eServer

The IBM brand name for a series of server products that are optimized for e-commerce. The products include the IBM System i, System p, IBM System x, and IBM System z.

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 System Storage Multipath Subsystem Device Driver (SDD)

Software that is designed to provide multipath configuration environment support for a host system that is attached to storage devices. SDD resides in a host system with the native disk device driver.

IBM System Storage Multipath Subsystem Device Driver Path Control Module (SDDPCM)

A loadable path control module for disk

storage system devices to supply path management functions and error recovery algorithms. When the disk storage system devices are configured as Multipath I/O (MPIO)-devices, SDDPCM becomes part of the AIX MPIO Fibre Channel Protocol device driver during the configuration. The AIX MPIO-capable device driver with the disk storage system SDDPCM module enhances the data availability and I/O load balancing.

IBM System Storage Subsystem Device Driver Device Specific Module (SDDDSM)

An IBM storage subsystems multipath I/O solution that is based on Microsoft MPIO technology. It is a device-specific module that is designed to support IBM storage subsystems devices such as SAN Volume Controller, DS8000, and DS6000. SDDDSM resides on a host server with the native disk device driver and provides enhanced data availability, automatic path failover protection, concurrent download of controller firmware code, and path selection policies for the host system.

IBM TotalStorage Enterprise Storage Server (ESS) A member of the Seascape product family of storage servers and attached storage devices (disk drive modules). The ESS provides for high-performance, fault-tolerant storage and management of enterprise data, providing access through multiple concurrent operating systems and communication protocols. High performance is provided by multiple symmetric multiprocessors, integrated caching, RAID support for the disk drive modules, and disk access through a high-speed serial storage architecture (SSA) interface.

IBM TotalStorage Enterprise Storage Server Expert (ESS Expert)

The software that gathers performance data from the ESS and presents it through a Web browser.

IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist)

Software with a Web-browser interface for configuring the ESS.

IBM TotalStorage Enterprise Storage Server Network (ESSNet)

A private network providing Web browser access to the ESS. IBM installs

the ESSNet software on an IBM workstation called the IBM TotalStorage ESS Master Console, supplied with the first ESS delivery.

IBM TotalStorage ESS Master Console (ESS Master Console)

An IBM workstation (formerly named the ESSNet console and hereafter referred to simply as the ESS Master Console) that IBM installs to provide the ESSNet facility when they install your ESS. It includes a Web browser that provides links to the ESS user interface, including ESS Specialist and ESS Copy Services.

ID See *identifier*.

identifier (ID)

A unique name or address that identifies things such as programs, devices, or systems.

IML See *initial microprogram 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. An initiator is typically an I/O adapter on a host computer. A SCSI initiator is analogous to an S/390 channel. A SCSI logical unit is analogous to an S/390 device. See also *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

In the ESS, 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 an ESA/390 command protocol). See 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)

In the ESS, an input-output adapter on the PCI bus.

IOCDS

See *input/output configuration data set*.

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

Facility provided by the Workload Manager of OS/390® and supported by

the ESS 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*.

IBM System i

An IBM eServer 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 CKD 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 use the key field to identify or locate a given record. The subsystem does not use the key field.

kilobyte (KB)

For processor storage, real, and virtual storage, and channel volume, 2¹⁰ or 1024 bytes.

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)

The algorithm used to identify and make available the cache space that contains the least-recently used data.

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*.

LIC See *Licensed Internal Code*.

Licensed Internal Code (LIC)

Microcode that IBM does not sell as part of a machine, but licenses to the customer. LIC 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. See also *licensed machine code (LMC)*.

LIFO See *last-in first-out*.

light-emitting diode (LED)

A semiconductor chip that gives off visible or infrared light when activated.

LMC See *licensed machine code*.

licensed machine code (LMC)

Microcode, basic input/output system code (BIOS), utility programs, device drivers, and diagnostics delivered with an IBM machine.

link address

On an ESCON or FICON interface, the portion of a source or destination address in a frame that ESCON or FICON uses to route a frame through an ESCON or FICON director. ESCON or FICON associates the link address with a specific switch port that is on the ESCON or FICON 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 or FICON hardware and logical functions of a control unit or channel subsystem that allow communication over an ESCON or FICON write interface and an ESCON or FICON 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 or FICON 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 ESS 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 ESS) 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 an S/390 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)

In Enterprise Systems Architecture/390, 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

In the ESS for Copy Services, a relationship between a source logical subsystem and target logical subsystem that is created over a physical path through the interconnection fabric used for Copy Services functions.

logical subsystem (LSS)

In the ESS, a topological construct that consists of a group of up to 256 logical devices. An ESS can have up to 16 CKD-formatted logical subsystems (4096 CKD logical devices) and also up to 16 fixed-block (FB) logical subsystems (4096 FB logical devices). The logical subsystem facilitates configuration of the ESS 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 an S/390 control-unit image.

For S/390 or IBM System z 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 used on a SCSI bus to enable it to differentiate between up to eight separate devices, each of which is a logical unit.

logical volume

The storage medium associated with a logical disk drive. A logical volume typically resides on one or more storage devices. The ESS administrator defines this unit of storage. The logical volume, when residing on a RAID 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 (DDM).

longitudinal redundancy check (LRC)

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.

In the ESS, a mechanism that the ESS uses for locating errors. The LRC checks the data as it progresses from the host, through the ESS controller, into the device adapter, and to the array.

longwave laser adapter

A connector used between a host and the ESS to support longwave fibre-channel communication.

loop The physical connection between a pair of device adapters in the ESS. 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.

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 Information Base (MIB)

A collection of objects that can be

accessed by means of a network management protocol. (GC)

In the ESS, 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*.

mass storage

The various techniques and devices for storing large amounts of data in a persisting and machine-readable fashion. Mass storage devices include all types of disk drives and tape drives.

master boot record

The boot program that the BIOS loads. This boot program is located in the first sector of the hard disk and is used to start the boot process

Master Console

See *IBM TotalStorage ESS Master Console*.

MB See *megabyte*.

MBR See *master boot record*.

MCA See *Micro Channel architecture*.

mean time between failures (MTBF)

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.

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 facility, the disk surface on which data is stored.

megabyte (MB)

For processor storage, real and virtual storage, and channel volume, 2^{20} or 1 048 576 bytes.

For disk storage capacity and communications volume, 1 000 000 bytes.

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

In the ESS, 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. When used in the context of data migration regarding the ESS, the transfer of data from one storage facility to another, such as from a 3390 to the ESS.

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.

mixed mode

With respect to a Copy Services domain, the mode of operation of the domain when one Copy Services server is a dual-active server and the other Copy Services server is a single-active server. See also *Copy Services server, dual-active server*, and *single-active server*.

MLC See *machine level control*.

mobile solutions terminal (MoST)

The mobile terminal used by service personnel.

mode conditioning patch

This cable is used to convert a single mode signal generated by 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 single mode light sent into a longwave adapter.

Model 100

A 2105 Model 100, often simply referred to as a Mod 100, is an expansion enclosure for the ESS. 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*.

multiple allegiance

An ESS hardware function that is independent of software support. This function enables multiple system images to concurrently access the same logical volume on the ESS as long as the system images are accessing different extents. See also *extent* and *parallel access volumes*.

multiple virtual storage (MVS)

Implies MVS/390, MVS/XA, MVS/ESA, and the MVS element of the OS/390 operating system.

multiplex

The action of transmitting simultaneously.

multiport serial adapter (MSA)

An adapter on the ESS Master Console that has multiple ports to which ESSs can be attached.

MVS See *multiple virtual storage*.

N

name server

A server that stores names of the participating ESS clusters.

Netfinity

IBM Intel-processor-based server; predecessor to the IBM xSeries server.

Netscape Navigator

Web browser software manufactured by Netscape.

Network Installation Management (NIM)

An environment that provides installation and configuration of software within a network interface.

NIM See *Network Installation Management*.

node

In a network, a point at which one or more functional units connect channels or data circuits. An ESS is a node in a fibre-channel network.

One SAN Volume Controller. Each node provides virtualization, cache, and Copy Services to the storage area network.

node failover

See *failover*.

non-RAID

A disk drive set up independently of other disk drives and not set up as part of a disk eight pack to store data using the redundant array of disks (RAID) data-stripping methodology.

nonremovable medium

A recording medium that cannot be added to or removed from a storage device.

nonvolatile storage (NVS)

In the ESS, 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 (IP) addressing, one of the four parts of a 32-bit integer presented in dotted decimal notation. dotted decimal notation consists of four 8-bit numbers written in base 10. For example, 9.113.76.250 is an IP address containing the octets 9, 113, 76, and 250.

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 ESS, such systems are those hosts that connect to the ESS through SCSI or FCP protocols. See also *small computer system interface* and *fibre-channel protocol*.

operating system (OS)

A set of programs that control how the system works. Controls the running of programs and provides such services as resource allocation, scheduling, input and output control, and data management.

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 See *operating system*.

OS/390

The IBM operating system that includes and integrates functions previously provided by many IBM software products (including the MVS operating system) for the IBM S/390 family of enterprise servers.

OS/400®

The IBM operating system that runs the IBM AS/400 and IBM System i eServer families of 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 ESS that enables OS/390 and z/OS systems to issue concurrent I/O requests against a CKD 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 ESA/390 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 ESA/390 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*.

PE See *IBM product engineering*.

Peer-to-Peer Remote Copy (PPRC)

A function 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 *synchronous PPRC* and *PPRC Extended Distance*.

peripheral component interconnect (PCI)

An architecture for a system bus and associated protocols that supports attachments of adapter cards to a system backplane.

persistent binding

A feature where a device has the same identification to the operating system after it restarts and after other devices are added to the operating system.

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 ESS (through cabling and switches) to a host adapter on another ESS.

point-to-point connection

For fibre-channel connections, a topology that enables the direct interconnection of ports. See *arbitrated loop* and *switched fabric*.

port In the ESS, a physical connection on a host adapter to the cable that connects the ESS to hosts, switches, or another ESS. The ESS 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.

PPRC See *Peer-to-Peer Remote Copy*.

PPRC Extended Distance

An optional feature for the ESS that maintains a fuzzy copy of a logical volume on the same ESS or on another ESS. 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 *Peer-to-Peer Remote Copy (PPRC)* and *synchronous PPRC*.

PPRC-XD

See *PPRC Extended Distance*.

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 Copy Services server

One of two Copy Services servers in a Copy Services server group. The primary Copy Services server is the active Copy Services server until it fails; it is then replaced by the backup Copy Services

server. A Copy Services server is software that runs in one of the two clusters of an ESS and performs data-copy operations within that group. See *active Copy Services server* and *backup Copy Services server*.

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 or bypass of a problem diagnosed by IBM in a current unaltered release of a program. (GC)

promote

To add a logical data unit to cache memory.

protected volume

In the IBM AS/400® platform, 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 ESS, however, can be configured to indicate that an AS/400 volume is protected or unprotected and give it RAID protection in either case.

System p

The product name of an IBM eServer product that emphasizes performance. It is the successor to the IBM RS/6000® family of servers.

pseudo-host

A host connection that is not explicitly defined to the ESS and that has access to at least one volume that is configured on the ESS. 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 ESS adds a pseudo-host icon only when the ESS 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 providing 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 *independent* 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 ESS 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 ESS 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 See *array*.

redundant array of independent disks (RAID)

A methodology of grouping disk drives for managing disk storage to insulate data from a failing disk drive.

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 acknowledgement and acknowledgement*.

request for acknowledgement and acknowledgement (REQ/ACK)

A cycle of communication between two data transport devices for the purpose of verifying the connection, which starts with a request for acknowledgement from one of the devices and ends with an acknowledgement 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

In Enterprise Systems Architecture/390, a relationship that is created in a control unit between a device and a channel path when the device completes a Sense Reserve command. The allegiance causes the control unit to guarantee access (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*.

RSSM IBM BladeCenter S SAS RAID Controller Module.

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 in the ESS as connected to S/390 servers. This term is synonymous with count-key-data (CKD) storage.

In ESS documentation, when noted, the term can refer to both S/390 and IBM System z storage. See also *IBM System z storage*.

SAID See *system adapter identification number*.

SAM See *sequential access method*.

SAN See *storage area network*.

SAS See *serial-attached SCSI*.

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 host systems

Host systems that are attached to the ESS with a SCSI interface. Such host systems run on UNIX, OS/400, Windows, 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.

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 ESS. See also *fibre-channel protocol* and *small computer system interface*.

SDD See *IBM System Storage Enterprise Storage Server Subsystem Device Driver*.

SDDDSM

See *IBM System Storage Subsystem Device Driver Device Specific Module*.

SDDPCM

See *IBM System Storage Multipath Subsystem Device Driver Path Control Module*.

Seascape architecture

A storage system architecture developed by IBM for open-systems servers and S/390 and IBM System z host systems. It provides storage solutions that integrate software, storage management, and technology for disk, tape, and optical storage.

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-attached SCSI (SAS)

A data transfer technology that uses a host bus adapter with four or more channels that operate simultaneously. Each full-duplex channel, known as a SAS port, transfers data in each direction.

serial connection

A method of device interconnection for determining interrupt priority by connecting the interrupt sources serially.

serial storage architecture (SSA)

An IBM standard for a computer peripheral interface. The interface uses a SCSI logical protocol over a serial

interface that configures attached targets and initiators in a ring topology. See also *SSA adapter*.

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 ESS, for example, in each host bay and in each cluster.

service information message (SIM)

A message sent by a storage server to service personnel through an S/390 operating system.

service personnel

A generalization referring to individuals or companies authorized to service the ESS. 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 used to service a storage facility.

service support representative (SSR)

Individuals or a company authorized to service the ESS. This term also refers to a service provider, a service representative, or an IBM service support representative (SSR). An IBM SSR installs the ESS.

session

A collection of multiple copy sets that comprise a consistency group.

shared storage

In an ESS, 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 used between host and ESS 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-active mode

With respect to a Copy Services server, the mode of operation of the server when the LIC level of the associated ESS cluster is below 2.0.

With respect to a Copy Services domain, the mode of operation of the domain when the Copy Services servers are single-active servers. See also *Copy Services server*, *Copy Services domain*, *dual-active server*, and *mixed mode*.

Single-Byte Command Code Sets Connection (SBCON)

The ANSI standard for the ESCON or FICON 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 *symmetric multiprocessor*.

SMS See *Systems Management Server*.

SNMP

See *simple network management protocol*.

Systems Management Server (SMS)

Change and configuration management software from Microsoft that runs on the

- Microsoft platform and distributes relevant software and updates to users.
- 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.
- spare** A disk drive on the ESS 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*.
- Specialist**
See *IBM TotalStorage Enterprise Storage Server Specialist*.
- Shared Product Object Tree (SPOT)**
A version of the /usr file system that diskless clients mount as their own /usr directory.

For NIM, a /usr file system or an equivalent file system that is exported by servers in the NIM environment for remote client use.
- SPOT** See *Shared Product Object Tree*.
- SSA** See *serial storage architecture*.
- SSA adapter**
A physical adapter based on serial storage architecture. SSA adapters connect disk drive modules to ESS clusters. See also *serial storage architecture*.
- SSID** See *subsystem identifier*.
- SSR** See *service support representative*.
- stacked status**
In Enterprise Systems Architecture/390, 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**
In the ESS, a volume that emulates one of several S/390 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 S/390 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 complex**
Multiple storage facilities.
- 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 facility**
A physical unit that consists of a storage server integrated with one or more storage devices to provide storage capability to a host computer.

A storage server and its attached storage devices.
- 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 are not provided by the storage device. The storage server has one or more clusters.
- striping**
A technique that distributes data in bit, byte, multibyte, record, or block increments across multiple disk drives.

subchannel

A logical function of a channel subsystem associated with the management of a single device.

Subsystem Device Driver

See *IBM System Storage Multipath Subsystem Device Driver*.

Subsystem Device Driver Device Specific Module (SDDDSM)

An IBM storage subsystems multipath I/O solution that is based on Microsoft MPIO technology. It is a device-specific module that is designed to support IBM storage subsystems devices such as SAN Volume Controller, DS8000, and DS6000. SDDDSM resides on a host server with the native disk device driver and provides enhanced data availability, automatic path failover protection, concurrent download of controller firmware code, and path selection policies for the host system.

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 ESS to receive a trace or dump package. See also *support catcher* and *remote technical assistance information network*.

switched fabric

In the ESS, one of three fibre-channel connection topologies that the ESS supports. See also *arbitrated loop* and *point-to-point*.

symmetric 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 ESS contains four processors set up in SMP mode.

synchronous PPRC

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 *Peer-to-Peer Remote Copy* and *PPRC Extended Distance*.

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)

In the ESS, the unique identification number automatically assigned to each ESS host adapter for use by ESS 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 slave 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 an S/390 control unit. A SCSI initiator is analogous to an S/390 channel. A SCSI logical unit is analogous to an S/390 device. See also *small computer system interface*.

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)

Nominally, 1 000 000 000 000 bytes, which is accurate when speaking of bandwidth and disk storage capacity.

For ESS cache memory, processor storage, real and virtual storage, a terabyte refers to 2⁴⁰ or 1 099 511 627 776 bytes.

terminal emulator

In the ESS, a function of the ESS Master 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.

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)

A combination of data-transmission protocols that provide end-to-end connections between applications over interconnected networks of different types.

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*.

U

UFS UNIX filing system.

Ultra-SCSI

An enhanced small computer system interface.

unconfigure

To delete the configuration.

unit address

In Enterprise Systems Architecture/390, the address associated with a device on a given control unit. On ESCON or FICON 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.

UNIX File System (UFS)

A section of the UNIX file tree that is physically contained on a single device or disk partition and that can be separately mounted, dismounted, and administered.

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 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). See ANSI X3.230-199x.

UTC See *Coordinated Universal Time*.

utility device

The ESA/390 term for the device used with the Extended Remote Copy facility to access information that describes the modifications performed on the primary copy.

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.

virtualization

In the storage industry, a concept in which a pool of storage is created that contains several disk subsystems. The subsystems can be from various vendors. The pool can be split into virtual disks that are visible to the host systems that use them. In SDD, virtualization product refers to SAN Volume Controller.

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/370, VM/ESA, VM/CMS, and VM/SP. See also *virtual machine (VM) facility*.

volume

In Enterprise Systems Architecture/390, 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.

VPD See *vital product data*.

VSE/ESA

IBM operating system, the letters of which represent virtual storage extended/enterprise systems architecture.

W**Web Copy Services**

See *ESS Copy Services*.

worldwide node name (WWNN)

A unique 64-bit identifier for a host containing 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.

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.

WWPN

See *worldwide port name*.

X

XD See *PPRC Extended Distance*.

XRC See *Extended Remote Copy*.

xSeries

The product name of an IBM eServer 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 eServer IBM System z family of servers uses the z/Architecture architecture. It is the successor to the S/390 and 9672 family of servers. See also *Enterprise Systems Architecture/390*.

z/OS An operating system for the IBM eServer product line that supports 64-bit real storage.

IBM System z

An IBM eServer family of servers that emphasizes near-zero downtime.

IBM enterprise servers based on z/Architecture.

IBM System z storage

Storage arrays and logical volumes that are defined in the ESS as connected to IBM System z servers. See also *S/390 storage*.

Index

Special characters

/opt/IBMsdd/bin/showvpath
command 338

A

about this book xiii

accessing

AIX

- add a data path volume to a volume group SMIT panel 85
- add a volume group with data path devices SMIT panel 84
- add paths to available data path devices SMIT panel 84
- backup a volume group with data path devices SMIT panel 85
- configure a defined data path device SMIT panel 84
- define and configure all data path devices SMIT panel 84
- display data path device adapter status SMIT panel 83
- display data path device configuration SMIT panel 82
- Display Data Path Device Status SMIT panel 83
- remake a volume group with data path devices SMIT panel 86
- remove a data path device SMIT panel 84
- Remove a Physical Volume from a Volume Group SMIT panel 85

adapters 299

configuring

Linux 204, 205

NetWare 299

Windows Server 2008 341

Windows Server 2012 341

firmware level 19, 108

LP7000E 14, 104

upgrading firmware level to (sf320A9) 109

adding

paths

AIX 46

Windows Server 2008 host systems 344

Windows Server 2012 host systems 344

addpaths

command 78, 87

utility programs, AIX 86

agreement for Licensed Internal Code 415

AIX

4.3.3

32-bit applications 20

64-bit applications 20

AIX (*continued*)

5.1.0

32-bit applications 20

64-bit applications 20

5.3.0

32-bit applications 20

64-bit applications 20

6.1.0

64-bit applications 20

accessing

add a data path volume to a volume group SMIT panel 85

add a volume group with data path devices SMIT panel 84

add paths to available data path devices SMIT panel 84

backup a volume group with data path devices SMIT panel 85

configure a defined data path device SMIT panel 84

define and configure all data path devices SMIT panel 84

display data path device adapter status SMIT panel 83

display data path device configuration SMIT panel 82

Display Data Path Device Status SMIT panel 83

remake a volume group with data path devices SMIT panel 86

Remove a Physical Volume from a Volume Group SMIT panel 85

removing a data path device SMIT panel 84

active/passive storage device

controller health check 126

applications

32-bit 20

64-bit 20

backing-up files belonging to an SDD volume group 80

changing SDDPCM controller healthcheck delay_time 133

changing the path-selection policy 69, 216

configuring

SDD 38, 44

volume group for failover protection 74

dynamically adding paths to SDD vpath devices of a volume

group 46

error messages 397

messages for persistent reserve environment 397

exporting

volume group with SDD 79

extending

an existing SDD volume group 80

fibre-channel

dynamic device tracking 70

AIX (*continued*)

fibre-channel (*continued*) requirements 16, 107

host attachment, installing 17

host system

disk driver 2

protocol stack 2

Hot Plug support 47

importing

volume group with SDD 78

installing

SDD 23

SDDPCM 111

migrating

an existing non-SDD volume group to SDD vpath devices in concurrent mode 95

non-SDD volume group to a SAN Volume ControllerViVo SDD multipath volume group in concurrent mode 93

non-SDD volume group to an ESS SDD multipath volume group in concurrent mode 93

MPIO devices 127

NIM SPOT server 114

OS

upgrade 31

recovering

from mixed volume groups 80

removing

SDD from a host system 49

SDDPCM from a host system 123

restoring files belonging to an SDD volume group 81

SDD utility programs 86

SDD-specific SMIT panels 81

supported fibre-channel device driver features 132

the loss of a device path 76

trace 97, 139

unconfiguring SDD 45

unsupported environments 13, 103

upgrading 28

verifying SDD

configuration 45

installation 35

AIX 5.1.0

32-bit 21

64-bit 21

AIX 5.2.0

32-bit 21

64-bit 21

AIX 5.3.0

32-bit 21

64-bit 21

automount, setting up 218

B

backing-up AIX files belonging to an SDD volume group 80
BladeCenter
 concurrent download of licensed machine code 8
block disk device interfaces (SDD) 177, 313
boot -r command 335
bootinfo -K command 21

C

cat /proc/modules command 209
cat /proc/scsi/scsi command 210
cat /proc/scsi/xxx/N command 210
cat /proc/sdd command 210
cd /media command 205, 206
cd /mnt command 205, 206
cd /opt/IBMsdd command 207
cd /opt/IBMsdd/bin command 208
cfallvpath 52
cfgmgr
 command 18, 78, 120
 run for each installed SCSI or fibre-channel adapter 78
 run n times where n represents the number of paths per SDD device. 78
cfgvpath command 213
changing
 path-selection policy for AIX 69, 216
 path-selection policy for HP 184
 path-selection policy for Solaris 324
 SDD hardware configuration
 HP-UX host systems 183
 Solaris host systems 321
 to the /dev directory
 HP-UX host systems 197
changing SDDPCM controller healthcheck delay_time 133
chdev command 77
chgrp command 335
chkconfig - -level X sdd off command 214
chkconfig - -level X sdd on command 214
chkconfig - -list sdd command 214
chkvpenv command 209
chmod command 335
command
 /opt/IBMsdd/bin/showvpath 338
 addpaths 78, 87
 boot -r 335
 bootinfo -K 21
 cat /proc/modules 209
 cat /proc/scsi/scsi 210
 cat /proc/scsi/xxx/N 210
 cat /proc/sdd 210
 cd /media 205, 206
 cd /mnt 205, 206
 cd /opt/IBMsdd 207
 cd /opt/IBMsdd/bin 208
 cfgmgr 18, 78, 120
 running n times for n-path configurations 78

command (continued)
 chdev 77
 chgrp 335
 chkconfig - -level X sdd off 214
 chkconfig - -level X sdd on 214
 chkconfig - -list sdd 214
 chkvpenv 209
 chmod 335
 datapath clear device count 372
 datapath disable ports 373
 datapath enable ports 374
 datapath open device path 374
 datapath query adapter 376
 datapath query adaptstats 377
 datapath query device 74, 78, 180, 213, 342, 378
 datapath query devstats 381
 datapath query essmap 383
 datapath query portmap 384
 datapath query version 385
 datapath query wwpn 386
 datapath remove adapter 386, 387
 datapath remove device 388
 datapath remove device path 387
 datapath set adapter 389
 datapath set device 0 path 0 offline 391
 datapath set device N policy rr/fo/lb/df 69, 185, 217, 325
 datapath set device path 390
 datapath set device policy 389
 datapath set performance 391
 datapath set qdepth 391, 392
 dpovgfix 73, 87
 dpovgfix vg-name 77
 esvpcfg 295
 excludesddcfg 51, 89
 extendvg 80
 extendvg4vp 80, 89
 hd2vp and vp2hd 87
 hd2vp vg_name 29
 HP-UX host systems
 hd2vp 184
 vgexport 193
 vgimport 194
 vp2hd 184
 vpcluster 194
 insmod ./vpath.o 209
 installp 17
 instfix -i | grep IY10201 17
 instfix -i | grep IY10994 17
 instfix -i | grep IY11245 17
 instfix -i | grep IY13736 17
 instfix -i | grep IY17902 17
 instfix -i | grep IY18070 17
 ls -al /unix 21
 ls -l 208
 lsfcg -vl fcsN 19, 108
 lsdev -Cc disk 18, 120
 lsdev -Cc disk | grep 2105 38
 lsdev -Cc disk | grep SAN Volume Controller 38
 lspp -l ibmSdd_432.rte 35
 lspp -l ibmSdd_433.rte 35, 36
 lspp -l ibmSdd_510.rte 35, 36
 lspp -l ibmSdd_510nchacmp.rte 35, 36

command (continued)
 lspv 28, 75
 lsvg -p vg-name 75
 lsvgfs 28
 lsvpcfg 28, 73, 77, 87, 213
 lsvpd 212
 metadb -a <device> 336
 metadb -d -f <device> 336
 metadb -i 336
 metainit 336
 metainit d <metadevice number> -t <"vpathNs" - master device> <"vpathNs" - logging device> 338
 metastat 336, 338
 mkdev -l vpathN 45
 mksysb restore command 76
 mkvg 74
 mkvg4vp 74, 88
 newfs 337
 odmget -q "name = ioaccess" CuAt 56
 orainst /m 332
 pcmpath chgprefercntl device 172, 173
 pcmpath clear device count 147
 pcmpath disable ports 147
 pcmpath enable ports 149
 pcmpath open device path 150
 pcmpath query adapter 152
 pcmpath query adaptstats 153
 pcmpath query device 154
 pcmpath query devstats 158
 pcmpath query essmap 160
 pcmpath query port 161
 pcmpath query portmap 162
 pcmpath query portstats 163
 pcmpath query session 165
 pcmpath query version 166
 pcmpath query wwpn 166
 pcmpath set adapter 166
 pcmpath set device 0 path 0 offline 171
 pcmpath set device algorithm 167
 pcmpath set device cntlhc_interval 170
 pcmpath set device health_check mode 169
 pcmpath set device path 171
 pcmpath set device svc_sb_ttl 172
 pcmpath set health_check time interval 168
 pkgmgr IBMsdd 337
 restvg 81
 restvg4vp 81
 rmdev 78
 rmdev -dl dpo -R 28, 49, 123
 rmdev -l dpo -R 45
 rpm -e IBMsdd command 217
 rpm -qi IBMsdd 207, 217
 rpm -ql IBMsdd 207, 217
 savevg 81
 savevg4vp 81
 showvpath 198, 333, 334, 336
 shutdown -i6 -y -g0 336
 shutdown -rF 18, 120
 smitty 28
 smitty deinstall 17

- command (*continued*)
 - smitty device 28
 - smitty uninstall 17
 - table of, in installation package 21
 - umount 28, 338
 - umount /cdrom 320
 - unmod ./sdd-mod.o 217
 - using 174
 - varyoffvg 28, 38
 - varyonvg vg_name 29
- comments, how to send xx
- concurrent download of licensed machine code
 - BladeCenter S SAS RAID Controller Module 8
 - disk storage systems 7
 - DS4000 7
 - RSSM 8
 - virtualization products 7
- configuring
 - AIX
 - disk storage system 15
 - ESS 106
 - fibre-channel attached devices 16, 107
 - fibre-channel-attached devices 18, 120
 - SAN Volume ControllerViVo 15
 - volume group for failover protection 74
 - clusters with SDD
 - Windows Server 2008 host systems 353
 - Windows Server 2012 host systems 353
 - disk storage system
 - NetWare host systems 299
 - DS4000
 - for HP-UX 176
 - ESS
 - HP-UX host systems 176
 - Linux host systems 204
 - Solaris host systems 312
 - fibre-channel adapters
 - Linux host systems 204, 205
 - NetWare host systems 299
 - Windows Server 2008 host systems 341
 - Windows Server 2012 host systems 341
 - MPIO devices as the SAN boot device 135
 - SAN Volume ControllerViVo
 - Solaris host systems 312
 - SDD
 - AIX 44
 - AIX host 38
 - at system startup 214
 - Linux host systems 208, 209
 - NetWare host systems 301
 - Solaris host systems 321
 - SDDDSM
 - Windows Server 2008 344
 - supported storage device
 - Windows Server 2008 341
 - Windows Server 2012 341
 - virtualization products 176

- configuring (*continued*)
 - Linux host systems 204
- controller health check feature, SDDPCM
 - active/passive storage device 126
- conversion scripts
 - hd2vp 86
 - vp2hd 45, 86
- creating
 - device node for the logical volume
 - device in an HP-UX host systems 197
 - directory in /dev for the volume group in an HP-UX host systems 197
 - file system on the volume group in an HP-UX host systems 198
 - logical volume in an HP-UX host systems 198
 - new logical volumes in an HP-UX host systems 197
 - physical volume in an HP-UX host systems 198
 - volume group in an HP-UX host systems 198
- customizing
 - Network File System file server 201
 - Oracle 331
 - standard UNIX applications 196, 329

D

- database managers (DBMS) 313
- datapath
 - clear device count command 372
 - commands 371
 - disable ports command 373
 - enable ports command 374
 - open device path command 374
 - query
 - adapter 387, 388
 - adapter command 376
 - adapter command changes 186
 - adaptstats command 377
 - device command 74, 78, 213, 378
 - devstats command 381
 - essmap command 383
 - portmap command 384
 - set adapter command 389
 - version command 385
 - wwpn command 386
 - remove
 - adapter 387
 - adapter command 386
 - device 388
 - device path command 387
 - set device 0 path 0 offline command 391
 - set device N policy rr/fo/lb/df command 69, 185, 217, 325
 - set device path command 390
 - set device policy command 389
 - set performance command 391
 - set qdepth 392
 - set qdepth command 391
- definitions 419

- determining
 - AIX
 - adapter firmware level 19, 108
 - major number of the logical volume device for an HP-UX host systems 197
 - size of the logical volume for an HP-UX host systems 199
 - device driver 312
 - devices.fcp.disk.ibm.rte 12, 16
 - devices.fcp.disk.ibm2105.rte 12
 - devices.scsi.disk.ibm2105.rte 12
 - disk storage system
 - configuring for NetWare 299
 - displaying
 - AIX
 - ESS SDD vpath device configuration 73
 - current version of SDD
 - Windows Server 2008 343
 - Windows Server 2012 343
 - dpovgfix command 73, 87
 - dpovgfix vg-name command 77
 - DS4000
 - concurrent download of licensed machine code 7
 - configuring LUNs attached to the HP-UX host system 176
 - installing an additional package for support 183
 - dynamic I/O load balancing 6
 - dynamically opening an invalid or close_dead path 70
 - dynamically removing or replacing adapters
 - AIX Hot Plug support 47
 - paths
 - AIX Hot Plug support 47
 - dynamically removing paths 48
 - dynamically replacing adapters
 - different type replacement 47
 - same type replacement 47
- E
 - enhanced data availability 4
 - error messages
 - AIX
 - messages for persistent reserve environment 397
 - VPATH_DEVICE_OFFLINE 397
 - VPATH_DEVICE_ONLINE 397
 - VPATH_PATH_OPEN 397
 - VPATH_XBUF_NOMEM 397
 - for ibmSdd_433.rte installation package for SDD
 - VPATH_FAIL_RELPRESERVE 397
 - VPATH_OUT_SERVICE 397
 - VPATH_RESV_CFLICT 398
 - SDDPCM 398
 - Windows 400, 401, 402, 403, 404
 - ESS
 - AIX
 - displaying SDD vpath device configuration 73
 - configuring for HP 176
 - configuring for Linux 204

- ESS (*continued*)
 - configuring for Solaris 312
 - devices (hdisks) 92
 - LUNs 92
- exporting a volume group with SDD, AIX 79
- extending an existing SDD volume group, AIX 80
- extendvg command 80
- extendvg4vp command 80, 89

F

- failover protection 6
 - AIX
 - creating a volume group from a single-path SDD vpath device 76
 - losing 75
 - manually deleted devices and running the configuration manager 78
 - side effect of running the disk change method 76
 - the loss of a device path 76
 - verifying load-balancing and failover protection 73
 - when it does not exist 73
- fibre-channel adapters
 - configuring
 - for Windows Server 2008 341
 - for Windows Server 2012 341
 - Linux host systems 204, 205
 - NetWare host systems 299
 - supported
 - AIX host systems 14, 104
 - HP-UX host systems 175
 - Linux host systems 203
 - NetWare host systems 299
 - Solaris host systems 311
 - Windows Server 2008 host systems 340
 - Windows Server 2012 host systems 340
- fibre-channel device drivers
 - configuring for AIX 16, 107
 - devices.common.IBM.fc 16, 107
 - devices.fcp.disk 16, 107
 - devices.pci.df1000f7 16, 107
 - installing for AIX 16, 107
 - supported
 - AIX host systems 14, 104
 - NetWare host systems 299
- fibre-channel requirements
 - Windows Server 2008 340
 - Windows Server 2012 340

H

- HACMP
 - concurrent mode 54
 - enhanced concurrent mode 124
 - hd2vp conversion script 56
 - importing volume groups 56
 - node fallover 66
 - nonconcurrent mode 54

- HACMP (*continued*)
 - persistent reserve 56
 - recovering paths 66
 - SDD persistent reserve attributes 56
 - software support for enhanced concurrent mode 124
 - software support for nonconcurrent mode 55
- hardware configuration
 - changing
 - HP-UX host systems 183
 - Solaris host systems 321
- hardware requirements
 - HP-UX host systems 175
 - Linux host systems 203
 - Solaris host systems 311
- hd2vp and vp2hd command 87
- hd2vp command
 - HP-UX host systems 184
- hd2vp vg_name command 29
- hdisk device
 - chdev 77
 - modify attributes 77
- healthcheck 130
- High Availability Cluster Multiprocessing (HACMP) 53
- host attachment
 - upgrade 31
- host system requirements
 - Windows Server 2008 340
 - Windows Server 2012 340
- HP-UX 11.0
 - 32-bit 176
 - 64-bit 176, 178
- HP-UX 11.11
 - 32-bit 178
 - 64-bit 178
- HP-UX 11.23
 - IA 64-bit 178
 - PA_RISC 64-bit 178
- HP-UX 11i
 - 32-bit 175
 - 64-bit 175
- HP-UX host systems
 - changing
 - SDD hardware configuration 183
 - the path-selection policy 184
 - to the /dev directory 197
 - creating
 - a file system on the volume group 198
 - a logical volume 198
 - a volume group 198
 - device node for the logical volume device 197
 - directory in /dev for the volume group 197
 - new logical volumes 197
 - physical volume 198
- determining
 - major number of the logical volume 197
 - size of the logical volume 199
- disk device drivers 188, 196
- disk driver 2
- dynamic reconfiguration 184
- installing SDD 179, 180

- HP-UX host systems (*continued*)
 - on a Network File System file server 201
 - on a system that already has NFS file server 202
- LJFS file system 201
- mounting the logical volume 198
- operating system 175
- protocol stack 2
- re-creating
 - existing logical volume 199
 - logical volume 200
 - physical volume 198, 200
 - volume group 200
- removing
 - existing logical volume 199
 - existing volume group 199
 - logical volumes 199
- SCSI disk driver (sdisk) 176
- SDD
 - installing 175
 - setting the correct timeout value for the logical volume manager 201
 - setting up Network File System for the first time 201
 - standard UNIX applications 196
 - support for DS4000 183
 - understanding how SDD works 176
 - unsupported environments 176
 - upgrading SDD 177, 181
 - using applications with SDD 196

I

- IBM System p with static LPARs
 - configured 20
- ibm2105.rte 15
- ibmSdd_433.rte installation package
 - for SDD 1.2.2.0
 - removing 56
 - for SDD 1.3.2.0. SDD vpath devices
 - unconfiguring 56
- importing a volume group with SDD, AIX 78
- insmod ./sdd-mod.o command 209
- installation package, AIX
 - devices.sdd.43.rte 20, 45
 - devices.sdd.51.rte 20, 45
 - devices.sdd.52.rte 20
 - devices.sdd.53.rte 20
 - devices.sdd.61.rte 20
 - devices.sdd.nn.rte 23
 - devices.sddpcm.52.rte 124
 - ibmSdd_432.rte 35, 50, 54, 93
 - ibmSdd_433.rte 35, 36, 50, 54, 56, 93, 397, 398
 - ibmSdd_510.rte 36, 50, 54
 - ibmSdd_510nchacmp.rte 50, 54
 - SDD 23
- installing
 - AIX
 - fibre-channel device drivers 16, 107
 - host attachment 17
 - planning 11, 99
 - SDD 23
 - SDDPCM 111

- installing (*continued*)
 - converting an Oracle installation from
 - sdisk on a Solaris host system 334
 - NetWare
 - planning 297
 - Oracle
 - Solaris host systems 331
 - package for DS4000 support 183
 - SDD
 - HP-UX host systems 175, 179, 180
 - Linux host systems 203, 205
 - NetWare host systems 301
 - on a Network File System file
 - server on a Solaris host
 - system 329
 - server on HP-UX host
 - systems 201
 - on a system that already has
 - Network File System file
 - server 330
 - on a system that already has NFS
 - file server on HP-UX host
 - systems 202
 - on a system that already has
 - Oracle on a Solaris host
 - system 333
 - on a system that already has
 - Solstice DiskSuite in place on a
 - Solaris host system 336
 - Solaris host systems 311, 315
 - using a file system on a Solaris
 - host system 333
 - using raw partitions on a Solaris
 - host system 334
 - SDDDSM
 - Windows Server 2008 host
 - systems 339, 341
 - Windows Server 2012 host
 - systems 339, 341
 - Solaris Volume Manager for the first
 - time on a Solaris host system 335
 - vpath on a system that already has
 - UFS logging in place on a Solaris
 - host system 337
 - installp command 17
 - instfix -i | grep IY10201 command 17
 - instfix -i | grep IY10994 command 17
 - instfix -i | grep IY11245 command 17
 - instfix -i | grep IY13736 command 17
 - instfix -i | grep IY17902 command 17
 - instfix -i | grep IY18070 command 17
 - invalid or close-dead path
 - dynamically opening 70

K

KB 160, 383

L

- Licensed Internal Code
 - agreement 415
- lilo, using with SDD (remote boot) on
 - x86 289

- Linux host systems
 - booting over the SAN with an IBM
 - SDD 224
 - configuring
 - ESS 204
 - fibres-channel adapters 204, 205
 - SDD 208, 214
 - virtualization products 204
 - disk driver 2
 - installing SDD 203, 205
 - logical volume manager 221
 - maintaining SDD vpath device
 - configuration persistency 215
 - partitioning SDD vpath devices 294
 - preparing SDD installation 204
 - protocol stack 2
 - removing SDD 217
 - unsupported environments 204
 - upgrading SDD 206
 - using
 - SDD configuration 209
 - standard UNIX applications 294
 - verifying SDD installation 207
 - load-balancing, AIX 73
 - loading SDD
 - on Linux 208, 209
 - on NetWare 301
 - on Solaris 321
 - logical volume manager 221, 313
 - losing failover protection, AIX 75
 - ls -al /unix command 21
 - ls -l command 208
 - lscfg -vl fcsN command 19, 108
 - lsdev -Cc disk | grep 2105 command 38
 - lsdev -Cc disk | grep SAN Volume
 - Controller command 38
 - lsdev -Cc disk command 18, 120
 - lslpp -l *Sdd* command 35
 - lslpp -l ibmSdd_432.rte command 35
 - lslpp -l ibmSdd_433.rte command 35, 36
 - lslpp -l ibmSdd_510.rte command 35, 36
 - lslpp -l ibmSdd_510nchacmp.rte
 - command 35, 36
 - lspv command 28, 75
 - lsvg -p vg-name command 75
 - lsvgfs command 28
 - lsvpcfg 52
 - lsvpcfg command 28, 73, 77, 87, 213, 295
 - lsvpcfg utility programs, AIX 87
 - lsvpd command 212

M

- maintaining SDD vpath device
 - configuration persistency, for Linux host
 - systems 215
 - manual exclusion of disk storage system
 - devices from the SDD configuration 51
 - metadb -a <device> command 336
 - metadb -d -f <device> command 336
 - metadb -i command 336
 - metainit command 336
 - metainit d <metadevice number> -t
 - <"vpathNs" - master device>
 - <"vpathNs" - logging device>
 - command 338
 - metastat command 336, 338

migrating

- AIX
 - an existing non-SDD volume
 - group to SDD vpath devices in
 - concurrent mode 95
 - non-SDD volume group to a SAN
 - Volume ControllerViVo SDD
 - multipath volume group in
 - concurrent mode 93
 - non-SDD volume group to an ESS
 - SDD multipath volume group in
 - concurrent mode 93
 - mirroring logical volumes 94
 - mkdev -l vpathN command 45
 - mksysb restore command 76
 - mkvg command 74
 - mkvg4vp command 74, 88
 - mounting the logical volume, HP 198
 - multipath SAN boot support 135

N

- NetWare host systems
 - configuring
 - disk storage system 299
 - fibres-channel adapters 299
 - SDD 301
 - error logging 304
 - error reporting 304
 - example command output 306
 - installing SDD 301
 - preparing SDD installation 299
 - removing SDD 305
 - supported environments 298
 - unsupported environments 298
 - newfs command 337
 - NIM environment
 - special considerations 50
 - uninstalling SDD 50
 - NIM SPOT server 114
 - notices
 - Licensed Internal Code 415
 - notices statement 413

O

- odmget -q "name = ioaccess" CuAt
 - command 56
- Open HyperSwap
 - quiesce expire time 131
 - Replication 9
- orainst /m command 332

P

- partitioning SDD vpath devices, for
 - Linux host systems 294
- path-failover protection system 6
- path-selection policy
 - changing 69, 185, 217, 325
 - default (optimized) 303
 - failover only 69, 185, 216, 303, 324
 - load balancing 69, 185, 216, 303, 324
 - round robin 69, 185, 216, 303, 324
- pcmpath
 - chgprefercntl device 173

- removing *(continued)*
 - SDD *(continued)*
 - from an AIX host system 49
 - Linux host systems 217
 - NetWare host systems 305
 - Windows Server 2008 host systems 350
 - Windows Server 2012 host systems 350
 - SDDDSM in a two-node cluster environment 354
 - SDDPCM
 - from an AIX host 123
 - from an AIX host system 123
- replacing manually excluded devices in the SDD configuration 52
- requirements
 - hardware and software
 - HP 175
 - Linux host systems 203
 - Solaris host systems 311
 - hardware, AIX
 - disk storage systems 12
 - Fibre channel adapters and cables 12
 - Host system 12
 - SAN Volume ControllerViVo 12
 - SCSI adapters and cables 12
 - hardware, SDDPCM
 - fibre adapters and cables 103
 - Fibre adapters and cables 103
 - Host system 103
 - supported storage devices 103
 - hardware, Windows Server 2008
 - disk storage system 339
 - hardware, Windows Server 2012
 - disk storage system 339
 - host system, AIX 13
 - ESS 14
 - Fibre 14
 - SAN Volume ControllerViVo 14
 - SCSI 14
 - supported storage devices 104
 - host system, NetWare 297
 - disk storage system 298
 - fibre 299
 - SCSI 298
 - host system, SDDPCM 104
 - Fibre 104
 - software
 - AIX operating system 12
 - ibm2105.rte ESS package, AIX 12
 - SCSI and fibre-channel device drivers, AIX 12
 - SDDPCM, AIX 5.2 TL07 (or later), AIX 5.3 TL08 (or later), or AIX 6.1 TL02 operating system 103
 - Windows Server 2008 operating system 339
 - Windows Server 2012 operating system 339
- reserve policy
 - attribute 44
 - controlling 44

- restoring
 - AIX
 - files belonging to an SDD volume group 81
 - restvg command 81
 - restvg4vp command 81
 - rmdev -dl dpo -R command 28, 49, 123
 - rmdev -l dpo -R command 45
 - rmdev command 78
 - rmvpath xxx command 213
 - rpm -e IBMsdd command 217
 - rpm -qi IBMsdd command 207, 217
 - rpm -ql IBMsdd command 207, 217
 - RSSM
 - concurrent download of licensed machine code 8

S

- SAN Volume ControllerViVo
 - configuring for Solaris 312
 - preferred node path selection algorithm 8
 - Preferred Node path selection algorithm 185
- savevg command 81
- savevg4vp command 81
- SCSI adapter support
 - HP-UX host systems 175
 - NetWare host systems 298
 - Solaris host systems 311
- SCSI-3 Persistent Reserve command set 56
- SDD
 - adding paths
 - to vpath devices of a volume group 46
 - architecture 2
 - automount 218
 - configuration
 - checking 45
 - host attachment, removing 50
 - how it works on HP-UX host systems 176
 - how it works on Solaris host systems 312
 - installing
 - AIX 11
 - HP-UX host systems 175
 - Linux 203
 - Linux over the SAN 224
 - NetWare 297
 - scenarios 177
 - Solaris host systems 311, 315
 - introduction 2
 - mounting devices with automount 218
 - NIM environment 50
 - overview 2
 - post-installation
 - HP-UX host systems 188
 - postinstallation
 - Solaris host systems 317
 - removing
 - HP-UX host systems 190
 - NetWare host systems 305
 - server daemon 365

- SDD *(continued)*
 - AIX host systems 21, 67, 292
 - HP-UX host systems 191
 - Solaris host systems 327
 - starting manually 192
 - stopping 192
 - upgrading
 - automatically 24
 - HP-UX host systems 177
 - manually 28
 - userspace commands for reconfiguration 213
 - using applications
 - with SDD on HP Network File System file server 201
 - with SDD on HP-UX standard UNIX applications 196
 - with SDD on Linux standard UNIX applications 294
 - with SDD on Solaris Network File System file Server 329
 - with SDD on Solaris standard UNIX applications 329
 - with SDD on Solaris, Oracle 331
 - utility programs, AIX 86
 - verifying
 - configuration 45
 - vpath devices 92
 - website xiii
- SDDDSM
 - 2.4.5.0 features 356
 - configuring
 - Windows Server 2008 344
 - datapath command support 354
 - displaying the current version on Windows Server 2008 343
 - displaying the current version on Windows Server 2012 343
 - installing
 - Windows Server 2008 host systems 339, 341
 - Windows Server 2012 host systems 339, 341
 - server daemon
 - Windows Server 2008 host systems 355
 - upgrading
 - Windows Server 2008 343
 - Windows Server 2012 343
 - verifying additional paths to SDDDSM devices 345, 347
- SDDDSM 2.4.5.0
 - failover group limit enhancement 363
 - FC in Hyper-V support 361
 - iSCSI in Hyper-V support 359
 - iSCSI support 358
 - load balance policies 356
 - Path Performance Plug-in 361
- sdddsm tools 405
- sddpcm 365
- SDDPCM
 - installing
 - AIX 99
 - from AIX NIM SPOT server 114
 - maximum number of devices 117
 - path healthcheck time interval 131

- SDDPCM (*continued*)
 - pcmpath commands 145
 - server 120
 - server daemon 139, 365
 - updating package 115, 117
 - verifying if the server has started 139
- SDDPCM server daemon 367
- sddserver.rte
 - AIX host systems 19
- sddsrv
 - AIX host systems 21, 67, 292
 - for ESS Expert
 - AIX host systems 19
 - for SDDPCM 365
 - HP-UX host systems 191
 - port binding 366
 - Solaris host systems 327
 - trace 366
 - Windows Server 2008 host systems 355
- sddsrv.conf file 367
- secondary-system paging space, managing 72
- setting up
 - correct timeout value for the logical volume manager on HP-UX host systems 201
 - Network File System for the first time on HP-UX host systems 201
 - NFS for the first time on a Solaris host system 329
 - Oracle using a file system
 - Solaris host systems 331
 - Oracle using raw partitions
 - Solaris host systems 332
 - UFS logging on a new system on a Solaris host system 337
- showvpath command 198, 333, 334, 336
- shutdown -i6 -y -g0 command 336
- shutdown -rF command 18, 120
- SMIT
 - definition 23, 112
- smitty command 28
- smitty deinstall command 17
- smitty device command 28
- smitty uninstall command 17
- smitty, definition 23, 112
- software requirements
 - for SDD on HP 175
 - for SDD on Linux 203
 - for SDD on Solaris 311
- Solaris 10 Zone support 323
- Solaris host systems
 - changing
 - SDD hardware configuration 321
 - the path-selection policy 324
 - configuring
 - SDD 321
 - disk device drivers 313
 - installing
 - Oracle 331
 - Solaris Volume Manager for the first time 335
 - vpath on a system that already has UFS logging in place 337
 - installing SDD 315

- Solaris host systems (*continued*)
 - converting an Oracle installation from sdisk 334
 - Network File System file server 329
 - system that already has Network File System file server 330
 - system that already has Oracle 333
 - system that already has Solstice DiskSuite in place 336
 - using a file system 333
 - using raw partitions 334
- operating system
 - upgrading SDD 311
- Oracle 331
- SCSI disk driver 312
- sd devices 329
- SDD 311
- SDD postinstallation 317
- setting up
 - NFS for the first time 329
 - UFS logging on a new system 337
- Solaris 10 Zone support 323
- Solaris Volume Manager 335
- Solstice DiskSuite 335
- standard UNIX applications 329
- supported environments 312
- UFS file system 329
- understanding how SDD works 312
- unsupported environments 312
- upgrading SDD 320
- upgrading the Subsystem Device Driver on 313
- using applications with SDD 329, 330
- Subsystem device driver, see SDD 335
- Sun host systems
 - disk driver 2
 - protocol stack 2
- support for
 - Windows Server 2008 353
 - Windows Server 2012 353
- supported environments
 - NetWare host systems 298
 - Solaris 312
- supported storage device
 - configuring on
 - Windows Server 2008 341
 - Windows Server 2012 341
- synchronizing logical volumes 95
- System Management Interface Tool (SMIT) 112
 - definition 23, 112
 - using for configuring 38
 - using to access the Add a Data Path Volume to a Volume Group panel on AIX host 85
 - using to access the Add a Volume Group with Data Path Devices panel on AIX host 84
 - using to access the backup a Volume Group with Data Path Devices on AIX host 85

- System Management Interface Tool (SMIT) (*continued*)
 - using to access the Configure a Defined Data Path Device panel on AIX host 84
 - using to access the Define and Configure All Data Path Devices panel on AIX host 84
 - using to access the Display Data Path Device Adapter Status panel on AIX host 83
 - using to access the Display Data Path Device Configuration panel on AIX host 82
 - using to access the Display Data Path Device Status panel on AIX host 83
 - using to access the Remake a Volume Group with Data Path Devices on AIX host 86
 - using to access the Remove a Data Path Device panel on AIX host 84
 - using to access the Remove a Physical Volume from a Volume Group panel on AIX host 85
 - using to back up a volume group with SDD on AIX host 80
 - using to backup a volume group with SDD on AIX host 85
 - using to create a volume group with SDD on AIX host 74
 - using to display the SAN Volume ControllerViVo SDD vpath device configuration on AIX host 73
 - using to display the SDD vpath device configuration on AIX host 73
 - using to export a volume group with SDD on AIX host 79
 - using to extend an existing SDD volume group on AIX host 80
 - using to import a volume group with SDD on AIX host 78
 - using to remove SDD from AIX host 49
 - using to remove SDDPCM from AIX host 123
 - using to restore a volume group with SDD on AIX host 86
 - using to restore a volume group with SDD vpath devices on AIX host 81
 - using to unconfigure SDD devices on AIX host 45
 - using to verify SDD configuration on AIX host 45

T
trademarks 415

U
umount
/cdrom command 320
command 28, 338
unconfiguring SDD on AIX 45

- understanding
 - how SDD works for HP-UX host systems 176
 - how SDD works for Solaris host systems 312
- unmod ./sdd-mod.o command 217
- unsupported environments
 - AIX 13, 103
 - HP 176
 - HP-UX 176
 - Linux 204
 - NetWare host systems 298
 - Solaris 312
 - Windows Server 2008 340
 - Windows Server 2012 340
- updating SDD using a PTF 29
- upgrade
 - AIX OS 31
 - host attachment 31
 - SDD packages 31
- upgrading
 - AIX
 - adapter firmware level 109
 - manually 28
 - SDD
 - automatically 24
 - HP-UX host systems 181
 - Linux host systems 206
 - Solaris host systems 320
 - SDD, manually
 - for AIX 4.3.2 28
 - for AIX 4.3.3 28
 - for AIX 5.1.0 28
 - for AIX 5.2.0 28
 - SDDDSM
 - Windows Server 2008 host systems 343
 - Windows Server 2012 host systems 343
- using
 - AIX trace function 97
 - datapath commands 371
 - ESS devices
 - directly, AIX 92
 - through AIX LVM 93
 - HP-UX applications with SDD 196
 - Linux standard UNIX applications 294
 - pcmpath commands 174
 - SAN Volume ControllerViVo devices through AIX LVM 93
 - SDDPCM trace function, AIX 138
 - Solaris applications with SDD 329
- utility programs, AIX
 - addpaths 86
 - dpovgfix 87
 - extendvg4vp 89
 - hd2vp and vp2hd 87
 - lsvpcfg 87
 - mkvg4vp 88
 - using disk storage system devices through AIX LVM 93
 - using ESS devices directly 92
 - using SAN Volume ControllerViVo devices through AIX LVM 93
 - using the SDDPCM trace function 138

- utility programs, AIX (*continued*)
 - using the trace function 97
- utility programs, HP
 - hd2vp 184
 - vp2hd 184
 - vpcluster 194

V

- varyoffvg command 28, 38
- varyonvg vg_name command 29
- verifying
 - additional paths are installed correctly
 - Windows Server 2008 host systems 345
 - Windows Server 2012 host systems 345
 - additional paths are removed correctly
 - Windows Server 2008 host systems 347
 - Windows Server 2012 host systems 347
- AIX
 - configuring SDD 45
 - SDD installation 35
 - SDD installation
 - Linux host systems 207
- Veritas Volume Manager
 - Command Line Interface for Solaris website 331
 - System Administrator's Guide website 331
- vgexport command
 - HP-UX host systems 193
- vgimport command
 - HP-UX host systems 194
- virtualization products
 - configuring for Linux 204
- volume groups
 - AIX 74
 - mixed
 - dpovgfix vg-name 77
 - how to fix problem 77
- vp2hd command
 - HP-UX host systems 184
- vpcluster command
 - HP-UX host systems 194

W

- websites
 - AIX APARs
 - maintenance level fixes and microcode updates 14
 - technology level fixes and microcode updates 104
 - HP-UX documentation 199
 - information on the fibre-channel adapters that can be used on NetWare host systems 299
 - information on the fibre-channel adapters that can be used on your AIX host 14, 104
 - information on the SCSI adapters that can attach to your NetWare host 298

- websites (*continued*)
 - NetWare APARs, maintenance level fixes and microcode updates 297
 - SDD xiii
- Windows Server 2008 host systems
 - adding
 - paths to SDDDSM devices 344
 - configuring
 - cluster with SDD 353
 - fibre-channel adapters 341
 - supported storage device 341
 - displaying the current version of the SDD 343
 - fibre-channel requirements 340
 - host system requirements 340
 - installing
 - SDDDSM 341
 - removing SDD 350
 - SDDDSM 339
 - support for clustering 353
 - unsupported environments 340
 - upgrading
 - SDDDSM 343
 - verifying
 - additional paths to SDDDSM devices 345, 347
- Windows Server 2012 host systems
 - adding
 - paths to SDDDSM devices 344
 - configuring
 - cluster with SDD 353
 - fibre-channel adapters 341
 - supported storage device 341
 - displaying the current version of the SDD 343
 - fibre-channel requirements 340
 - host system requirements 340
 - installing
 - SDDDSM 341
 - removing SDD 350
 - SDDDSM 339
 - support for clustering 353
 - unsupported environments 340
 - upgrading
 - SDDDSM 343
 - verifying
 - additional paths to SDDDSM devices 345, 347



Printed in USA

GC52-1309-08

