IBM TotalStorage Enterprise Storage Server
IBM TotalStorage SAN Volume Controller
IBM TotalStorage SAN Volume Controller for Cisco MDS 9000

Subsystem Device Driver User's Guide

Read Before Using

The IBM License Agreement for Machine Code is included in this book. Carefully read the agreement. By using this product, you agree to abide by the terms of this agreement and applicable copyright laws.

SC26-7637-00

Note
Before using this information and the product it supports, read the information in “Notices” on page 309.

First Edition (May 2004)

This edition replaces SC26-7507-00, and it includes information that specifically applies to the Subsystem Device Driver (SDD) Version 1 Release 5 Modification 1 Level 0 and including:

- Version 1 Release 4 Modification x Level x (or later) for IBM AIX® 4.3.2, AIX 4.3.3, AIX 5.1.0, AIX 5.2.0
- Version 1 Release 4 Modification 0 Level 0 (or later) for HP-UX 11.0 and HP-UX11i
- Version 1 Release 4 Modification x Level x (or later) for Linux Red Hat 7.x, Red Hat AS2.1, SuSE SLES7, SuSE SLES8 / UnitedLinux 1.0 (for Intel® i686 and IBM eServer pSeries®), Red Hat EL3
- Version 1.0.0.g (or later) for Novell
- Version 1 Release 4 Modification 0 Level 0 (or later) for Solaris 2.6, Solaris 7, Solaris 8, Solaris 9
- Version 1 Release 4 Modification 0 Level 0 (or later) for Microsoft® Windows NT® 4.0 Service Pack 6A or later
- Version 1 Release 4 Modification 0 Level 0 (or later) for Windows® 2000 Service Pack 2 or later
- Version 1 Release 4 Modification x Level x (or later) for Windows Server 2003

This edition also applies to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright International Business Machines Corporation 1999, 2004. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures xiii

Tables xv

About this book xvii
Who should use this book xvii
Command syntax conventions xvii
Highlighting conventions xvii
Special characters conventions xviii
Summary of changes xviii
New information xviii
Modified information xix
Related information xx
The ESS library. xx
The SAN Volume Controller Library xxii
The SAN Volume Controller for Cisco MDS 9000 library xxii
The SAN File System library xxiii
Other IBM publications xxiv
Ordering IBM publications xxx
Non-IBM publications xxxi
How to send your comments xxxii

Chapter 1. Overview of SDD 1
The SDD architecture 1
Enhanced data availability 3
Dynamic I/O load balancing. 6
Automatic path-failover protection 6
Concurrent download of ESS licensed internal code. 7
Concurrent download of SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 licensed internal code 7

Chapter 2. Using SDD on an AIX host system 9
Verifying the hardware and software requirements 9
Hardware 9
Software. 9
Unsupported environments 10
Host system requirements. 10
Preparing for SDD installation 12
Configuring the ESS 13
Configuring the SAN Volume Controller 13
Configuring the SAN Volume Controller for Cisco MDS 9000 13
Installing the AIX fibre-channel device drivers. 14
Uninstalling the AIX fibre-channel device drivers. 15
Configuring fibre-channel-attached devices 15
Removing fibre-channel-attached devices 16
Verifying the adapter firmware level 16
Determining if the sddServer for Expert is installed. 16
Planning for SDD installation on a pSeries 690 server LPAR 17
Determining the installation package 17
Determining the installation type 20
Installing SDD 21
Verifying the currently installed version of SDD for SDD 1.3.3.11 (or earlier) 22
Verifying your currently installed version of SDD for SDD 1.4.0.0 (or later) 23

© Copyright IBM Corp. 1999, 2004

Maximum number of LUNs 24
Preparing to configure SDD 28
Controlling I/O flow to SDD devices with the SDD qdepth_enable attribute . . . 29
Configuring SDD 30
Unconfiguring SDD 31
Migrating or upgrading SDD packages automatically without system restart 32
Preconditions for migration or upgrade 32
Procedures for migration or upgrade 32
Error recovery for migration or upgrade 33
Migrating or upgrading SDD manually 33
Migrating or upgrading the SDD package during an AIX OS upgrade 35
Updating SDD packages by applying a program temporary fix 35
Committing or Rejecting a PTF Update 36
Verifying the SDD configuration 38
Removing SDD from an AIX host system 39
Preferred node path-selection algorithm for SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 40
Dynamically changing the SDD path-selection policy algorithm 40
datapath set device policy command 41
Dynamically adding paths to SDD devices of a volume group 41
Dynamically opening an invalid or close_dead path 43
Dynamically replacing adapters or removing paths 43
Dynamically replacing another adapter of the same type 43
Dynamically replacing an adapter with an adapter of a different type 44
Dynamically removing a path of a SDD vpath device 46
Fibre-channel Dynamic Device Tracking for AIX 5.20 ML1 46
Using ESS concurrent download of licensed internal code 47
SDD server daemon 48
Verifying if the SDD server has started 48
Starting the SDD server manually 48
Changing to a different port number for the SDD server 48
Stopping the SDD server 48
PTFs for APARs on AIX with Fibre Channel and the SDD server 49
Understanding SDD 1.3.2.9 (or later) support for single-path configuration for ESS 50
Understanding SDD error recovery policies 50
Enabling fast failover to reduce error recovery time 52
Understanding SDD support for pSeries 690 with static LPARs configured . . . 52
Understanding the persistent reserve issue when migrating from SDD to non-SDD volume groups after a system reboot 52
Understanding SDD support for High Availability Cluster Multi-Processing . . . 53
HACMP and sddsrv 57
SDD persistent reserve attributes 58
Special requirements. 58
HACMP RAID concurrent-mode volume groups and enhanced concurrent-capable volume groups 60
HACMP RAID concurrent-mode volume groups 60
Enhanced concurrent-capable volume groups 64
Recovering paths that are lost during HACMP node fallover 67
Supporting enhanced concurrent mode in an HACMP environment 67
Managing secondary-system paging space 68
Listing paging spaces 68
Adding a paging space 68
Removing a paging space 68
Providing load-balancing and failover protection 69
Displaying the supported storage device vpath device configuration 69

Configuring volume groups for failover protection 70
Losing failover protection 71
Through the loss of a device path 72
By creating a volume group from single-path vpath (pseudo) devices 72
As a side effect of running the disk change method 72
By manually deleting devices and running the configuration manager (cfgmgr) 74
Importing volume groups with SDD 75
Importing volume groups under HACMP 76
Exporting a volume group with SDD 76
Recovering from mixed volume groups 77
Extending an existing SDD volume group 77
Backing up all files belonging to an SDD volume group 78
Restoring all files belonging to an SDD volume group 78
SDD-specific SMIT panels 78
Accessing the Display Data Path Device Configuration SMIT panel 79
Accessing the Display Data Path Device Status SMIT panel 80
Accessing the Display Data Path Device Adapter Status SMIT panel 80
Accessing the Define and Configure All Data Path Devices SMIT panel . . . 81
Accessing the Add Paths to Available Data Path Devices SMIT panel 81
Accessing the Configure a Defined Data Path Device SMIT panel 81
Accessing the Remove a Data Path Device SMIT panel 82
Accessing the Add a Volume Group with Data Path Devices SMIT panel 82
Accessing the Add a Data Path Volume to a Volume Group SMIT panel . . . 82
Accessing the Remove a Physical Volume from a Volume Group SMIT panel 83
Accessing the backup a Volume Group with Data Path Devices SMIT panel 83
Accessing the Remake a Volume Group with Data Path Devices SMIT panel 84
SDD utility programs 85
addpaths 85
hd2vp and vp2hd 85
dpovgfix 85
lsvpcfg 86
mkvg4vp 87
extendvg4vp 87
querysn 87
Persistent reserve command tool 88
Using supported storage devices directly 90
Using supported storage devices through AIX LVM. 91
Migrating a non-SDD volume group to a supported storage device SDD multipath volume group in concurrent mode 91
Migrating an existing non-SDD volume group to SDD vpath devices in concurrent mode 93
Using the trace function 95
Error messages 96
Error messages for the HACMP environment 96
Error messages for AIX Hot Plug support 97

Chapter 3. Using SDDPCM on an AIX host system 99
Supported SDDPCM features 100
Unsupported SDDPCM features 100
Verifying the hardware and software requirements 101
Hardware 101
Software 101
Unsupported environments 101
Host system requirements 101
Preparing for SDDPCM installation 102
Preparing for SDDPCM installation on an ESS 102

Installing SDDPCM 105
Creating and mounting the CD-ROM filesystem 106
Using the System Management Interface Tool facility to install SDDPCM. 107
Unmounting the CD-ROM File System 108
Verifying the currently installed version of SDDPCM 108
Maximum number of devices supported by SDDPCM 108
Configuring and unconfiguring ESS MPIO-capable devices 108
Verifying the SDDPCM Configuration 109
Updating SDDPCM 109
Updating SDDPCM packages by applying a program temporary fix 110
Committing or rejecting a program temporary fix update 111
Configuring ESS MPIO-capable devices as the boot device 111
Migrating from SDDPCM to the AIX default PCM or to SDD 112
Support system dump device with the ESS MPIO-capable device 112
SDDPCM ODM attribute settings 112
SDDPCM ODM attribute default settings 112
Changing device reserve policies 112
Changing the path selection algorithm 113
Changing SDDPCM path healthcheck mode 113
Changing SDDPCM path healthcheck time interval 114
Dynamically enabling and disabling paths or adapters 114
Dynamically enabling or disabling a path 114
Dynamically enabling or disabling an adapter 115
Dynamically adding and removing paths or adapters 115
Removing SDDPCM from an AIX host system 116
Using SDDPCM pcmpath commands 117
pcmpath disable ports 119
pcmpath enable ports 121
pcmpath open device path 123
pcmpath query adapter 125
pcmpath query adaptstats 126
pcmpath query device 127
pcmpath query devstats 129
pcmpath query essmap 131
pcmpath query portmap 132
pcmpath query wwpn 133
pcmpath set adapter 134
pcmpath set device path 135
SDDPCM server daemon 136
Verifying if the SDDPCM server has started 136
Starting the SDDPCM server manually 136
Stopping the SDDPCM server 136
Using the SDDPCM trace function 137
AIX 5.2.0 ML3 fibre-channel device driver features 137
Fast I/O failure of fibre-channel devices 137
Fibre channel dynamic device tracking 138

Chapter 4. Using SDD on a HP-UX host system 141
Verifying the hardware and software requirements 141
Hardware 141
Software 141
Unsupported environments 141
SDD support for 32-bit and 64-bit applications on HP-UX 11.0 and HP-UX 11i (fibre-channel) 142
Understanding how SDD works on an HP-UX host system 142
Preparing for SDD installation 142

Configuring the ESS 142
Configuring the SAN Volume Controller 142
Configuring the SAN Volume Controller for Cisco MDS 9000 142
Planning for installation 142
Determining if the SDD 1.3.1.5 (or later) server for Expert is installed . . . 144
Installing SDD 144
Postinstallation 146
Upgrading SDD 149
Upgrading SDD 1.3.0.2 (or earlier) to SDD 1.5.0.4 (or later) 149
Uninstalling the Subsystem Device Driver 150
Configuring SDD 150
Changing an SDD hardware configuration 151
Converting a volume group 151
Dynamic reconfiguration 151
Dynamically changing the SDD path-selection policy algorithm 152
Preferred node path-selection algorithm for the SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 152
SDD datapath query adapter command changes for SDD 1.4.0.0 (or later) 153
SDD datapath query device command changes for SDD 1.4.0.0 (or later) 154
SDD server daemon 155
Verifying if the SDD server has started 155
Starting the SDD server manually 155
Changing to a different port number for the SDD server 155
Stopping the SDD server 155
Understanding the SDD 1.3.1.5 (or later) support for single-path configuration for ESS 155
Understanding the SDD error recovery policy 156
Single-path mode (for ESS only) 156
Multipath mode 157
Dynamically opening an Invalid or CLOSE_DEAD path 157
How to import and export volume groups 157
Exporting volume groups 158
Moving the map file 158
Creating the volume group device directory 159
Creating the group special file 159
Importing volume groups 159
Using applications with SDD 160
Standard UNIX applications 160
Installing SDD on a Network File System file server 165
Oracle 166

Chapter 5. Using SDD on a Linux host system 171
Verifying the hardware and software requirements 171
Hardware 171
Software 171
Unsupported environments 172
Preparing for SDD installation 172
Configuring the ESS 172
Configuring the SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 172
Configuring the SAN File System 173
Configuring fibre-channel adapters for ESS 173
Configuring fibre-channel adapters for SAN Volume Controller 173
Configuring fibre-channel adapters for SAN Volume Controller for Cisco MDS 9000 173
Installing SDD 174

Upgrading the SDD 174
Verifying the SDD installation 175
Configuring SDD 176
Using multiple commands for SDD configuration 176
Using a single command for SDD configuration 180
Configuring SDD at system startup 180
Maintaining vpath device configuration persistence 181
Dynamically changing the SDD path-selection policy algorithm 181
Dynamically adding paths to SDD devices 182
Dynamic reconfiguration 182
Removing SDD 183
SDD server daemon 183
Verifying if the SDD server has started. 183
Starting the SDD server manually 183
Changing to a different port number for the SDD server 184
Stopping the SDD server. 184
Understanding the SDD error recovery policy 184
Understanding the SDD support for single-path configuration 185
Partitioning SDD vpath devices 185
Using standard UNIX applications 186

Chapter 6. Using SDD on a NetWare host system 187
Verifying the hardware and software requirements 187
Hardware requirements 187
Software requirements 187
Supported environments 188
Unsupported environments 188
ESS requirements 188
SCSI requirements 188
Fibre-channel requirements 188
Preparing for SDD installation 189
Configuring the ESS 189
Configuring fibre-channel adapters 189
Configuring SCSI adapters 190
Using a NetWare Compaq Server 190
Installing SDD 191
Maximum number of LUNs 191
Configuring SDD 191
Displaying the current version of the SDD 191
Features 191
Automatic path detection, failover and selection 191
Manual operations using the datapath commands 192
Understanding SDD error recovery algorithms 193
Dynamic load balancing 194
ESS logical unit detection 194
Error reporting and logging 194
SDD in NetWare-layered architecture 194
Removing SDD 195
Cluster setup for Novell NetWare 5.1 195
Cluster setup for Novel NetWare 6.0 195
Examples of commands output on the Console Window 195

Chapter 7. Using SDD on a Solaris host system 201
Verifying the hardware and software requirements 201
Hardware 201
Software 201

Supported environments 201
Unsupported environments 202
Understanding how SDD works on a Solaris host system 202
Preparing for SDD installation 202
Configuring the ESS 202
Configuring the SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 202
Determining if the SDD server for Expert is installed. 202
Planning for installation 203
Installing SDD. 204
Verifying the SDD installation 207
Uninstalling SDD. 207
Postinstallation 208
Upgrading the SDD. 210
Configuring SDD. 210
Changing an SDD hardware configuration 210
Dynamically changing the SDD path-selection policy algorithm 212
Preferred node path-selection algorithm for the SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000. 213
Understanding the SDD 1.3.1.0 (or later) support for single-path configuration for ESS 213
Understanding the SDD error recovery policy 214
SDD server daemon 216
Verifying if the SDD server has started. 216
Starting the SDD server manually 216
Changing to a different port number for the SDD server 216
Stopping the SDD server. 216
Using applications with SDD 216
Standard UNIX applications. 217
Installing SDD on a Network File System file server 217
Veritas Volume Manager 218
Oracle 219
Solaris Volume Manager (formerly Solstice DiskSuite) 222

Chapter 8. Using SDD on a Windows NT host system 227
Verifying the hardware and software requirements 227
Hardware 227
Software. 227
Unsupported environments 227
ESS requirements 228
Host system requirements 228
Preparing for SDD installation 229
Configuring the ESS 229
Configuring the SAN Volume Controller for Cisco MDS 9000 229
Configuring fibre-channel adapters 229
Configuring SCSI adapters for ESS devices. 229
Installing SDD. 230
Configuring SDD. 230
Adding paths to SDD devices 231
Preferred node path-selection algorithm for the SAN Volume Controller for Cisco MDS 9000 234
Upgrading the SDD. 235
Adding or modifying a multipath storage configuration to the supported storage device. 236
Removing SDD 238
SDD server daemon 238

Verifying that the SDD server has started . . . 238
Starting the SDD server manually . . . 238
Changing to a different port number for the SDD server . . . 239
Stopping the SDD server . . . 239
Displaying the current version of SDD . . . 239
Error recovery and retry policy . . . 239
Using high-availability clustering on an ESS . . . 241
Special considerations in the high-availability clustering environment . . . 241
Configuring a Windows NT cluster with SDD installed . . . 241

Chapter 9. Using SDD on a Windows 2000 host system . . . 245
Verifying the hardware and software requirements . . . 245
Unsupported environments . . . 245
ESS requirements . . . 245
SAN Volume Controller requirements . . . 245
SAN Volume Controller for Cisco MDS 9000 requirements . . . 246
Host system requirements . . . 246
Preparing for SDD 1.5.0.0 (or later) installation . . . 247
Configuring the supported storage device . . . 247
Configuring fibre-channel adapters . . . 247
Configuring SCSI adapters for ESS devices . . . 248
Installing SDD 1.5.0.0 (or later) . . . 248
Displaying the current version of SDD . . . 250
Upgrading SDD . . . 250
Upgrading to SDD 1.4.0.0 (or later) in a two-node cluster environment 251
Remote boot support for ESS . . . 251
Booting from an ESS device with Windows 2000 and SDD 1.5.0.0 (or later) using QLOGIC HBA . . . 251
Limitations when booting from an ESS device on a Windows 2000 host 252
Uninstalling SDD . . . 252
Removing SDD in a two-node cluster environment . . . 253
SDD server daemon . . . 253
Verifying if the SDD server has started . . . 253
Starting the SDD server manually . . . 253
Changing to a different port number for the SDD server . . . 254
Stopping the SDD server . . . 254
Adding paths to SDD devices . . . 254
Activating additional paths . . . 255
Verifying that additional paths are installed correctly . . . 255
Preferred Node path-selection algorithm for the SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 . . . 257
Error recovery and retry policy . . . 257
Support for Windows 2000 clustering . . . 259
Special considerations in the Windows 2000 clustering environment . . . 259
Configuring a Windows 2000 cluster with SDD installed . . . 259

Chapter 10. Using SDD on a Windows Server 2003 host system . . . 263
Verifying the hardware and software requirements . . . 263
Unsupported environments . . . 263
ESS requirements . . . 263
Host system requirements . . . 264
Preparing for SDD 1.5.0.0 (or later) installation . . . 264
Configuring the ESS . . . 265
Configuring the SAN Volume Controller . . . 265
Configuring fibre-channel adapters . . . 265
Configuring SCSI adapters . . . 265

Installing SDD 1.5.0.0 (or later) 266
Displaying the current version of SDD 268
Upgrading SDD 268
Upgrading from a Windows NT host system to Windows Server 2003 . . . 268
Uninstalling SDD 269
Removing SDD in a two-node cluster environment 269
SDD server daemon 270
Verifying if the SDD server has started 270
Starting the SDD server manually 270
Changing to a different port number for the SDD server 270
Stopping the SDD server 270
Adding paths to SDD devices 270
Activating additional paths 271
Verifying that additional paths are installed correctly 272
Error recovery and retry policy 273
Support for Windows Server 2003 clustering 274
Special considerations in the Windows Server 2003 clustering environment 275
Configure Windows 2003 cluster with SDD installed 275

Chapter 11. Using the SDD server and the SDDPCM server 279
SDD server daemon 279
Understanding how the SDD server daemon works 279
SDDPCM server daemon 280
sddsrv.conf and pcmsrv.conf file format 281
Enabling or disabling the sddsrv or pcmsrv TCP/IP port 281
Changing the sddsrv or pcmsrv TCP/IP port number 282

Chapter 12. Using the datapath commands 283
datapath disable ports 284
datapath enable ports 285
datapath open device path 286
datapath query adapter 288
datapath query adaptstats 290
datapath query device 291
datapath query devstats 293
datapath query essmap 295
datapath query portmap 296
datapath query wwpn 297
datapath remove adapter 298
datapath remove device path 299
datapath set adapter 301
datapath set device policy 303
datapath set device path 304
datapath set qdepth 305

Appendix. SDD data collection for problem analysis 307

Notices 309
Trademarks 310
IBM agreement for licensed internal code 311
Actions you must not take 312

Glossary 313

Index 333

Figures

1. Multipath connections between a host system and the disk storage in an ESS 4
2. Multipath connections between a host system and the disk storage with a SAN Volume Controller 5
3. Multipath connections between a host system and the disk storage with a SAN Volume Controller for Cisco MDS 9000 6
4. IBMsdd Driver 32-bit 145
5. IBMsdd Driver 64-bit 145
6. Example showing ESS devices to the host and path access to the ESS devices in a successful SDD installation on a Windows 2000 host system 250
7. Example showing ESS devices to the host and path access to the ESS devices in a successful SDD installation on a Windows Server 2003 host system 267

© Copyright IBM Corp. 1999, 2004

Tables

1. Publications in the SAN Volume Controller library xxii
2. Publications in the SAN Volume Controller for Cisco MDS 9000 library xxiii
3. Publications in the SAN File System library. xxiii
4. Other IBM publications xxiv
5. SDD in the protocol stack 1
6. Package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0 (or later) 10
7. SDD 1.4.0.0 (or later) installation packages for 32-bit and 64-bit applications on AIX 4.3.3 or later 17
8. Major files included in the SDD installation package 19
9. Maximum LUNs allowed for different AIX OS levels and different types of devices 25
10. Maximum SDD device configuration for ESS LUNs on AIX 5.2 25
11. List of previously installed installation packages that are supported with the installation upgrade 34
12. PTFs for APARs on AIX with fibre-channel support and the SDD server daemon running 49
13. Recommended SDD installation packages and supported HACMP modes. 53
14. Software support for HACMP 4.4, 4.4.1, and 4.5 on AIX 4.3.3 (32-bit only), 5.1.0 (32-bit and 64-bit), and 5.2.0 (32-bit and 64-bit) 54
15. Software support for HACMP 4.4, 4.4.4, and 4.5 on AIX 5.1.0 (32-bit and 64-bit kernel) 55
16. Software support for HACMP 4.3.1 and 4.4 in concurrent mode on AIX 4.3.3. 56
17. Software support for HACMP 4.3.1 and 4.4 in nonconcurrent mode on AIX 4.3.3 57
18. APARs on AIX with HACMP and the SDD server daemon running. 57
19. SDD-specific SMIT panels and how to proceed. 78
20. Commands 117
21. SDD installation scenarios 143
22. Patches necessary for proper operation of SDD on HP-UX 11.0 143
23. SDD components installed for HP-UX host systems 146
24. System files updated for HP-UX host systems 147
25. SDD commands and their descriptions for HP-UX host systems 147
26. SDD components for a Linux host system 175
27. Summary of SDD commands for a Linux host system. 176
28. SDD installation scenarios 203
29. Operating systems and SDD package file names 204
30. SDD components installed for Solaris host systems 208
31. System files updated for Solaris host systems. 209
32. SDD commands and their descriptions for Solaris host systems 209
33. Commands 283

© Copyright IBM Corp. 1999, 2004

About this book

The IBM Subsystem Device Driver (SDD) provides the multipath configuration environment support for a host system that is attached to:

- IBM TotalStorage® Enterprise Storage Server® (ESS)
- IBM TotalStorage SAN Volume Controller
- IBM TotalStorage SAN Volume Controller for Cisco MDS 9000

as well as a host system that is using an IBM TotalStorage SAN File System.

In this book, the phrase *supported storage devices* will be used to refer to the following types of devices:

- ESS
- ESS or SAN Volume Controller or both
- ESS or SAN Volume Controller for Cisco MDS 9000 or both

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

- IBM AIX®
- HP-UX
- Linux Red Hat 7.x and Advanced Server 2.1 (Intel®), SuSE SLES 7 (Intel), SuSE SLES8 / UnitedLinux 1.0 (for Intel i686 and IBM @server pSeries®), and Enterprise Linux 3.
- Netware (ESS only)
- Solaris
- Microsoft® Windows NT®
- Microsoft Windows® 2000
- Microsoft Windows Server 2003

Who should use this book

This book is intended for storage administrators, system programmers, and performance and capacity analysts.

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`).

© Copyright IBM Corp. 1999, 2004

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 **`lslpp -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.

Summary of changes

This book contains both information previously presented in the First Edition (November 2003) of the *IBM TotalStorage Enterprise Storage Server: Subsystem Device Driver User's Guide* and major technical changes to that information.

Note: For the last-minute changes that are not included in this book, see the README file on the SDD compact disc or visit the SDD Web site at:

www-1.ibm.com/servers/storage/support/software/sdd.html

New information

This edition includes the following new information:

- AIX:
  - Addition of AIX Hot Plug support.
- SDDPCM:
  - Support for the following new pcmpath commands:
    - **pcmpath disable ports**
    - **pcmpath enable ports**
    - **pcmpath open device path**
    - **pcmpath query essmap**
    - **pcmpath query portmap**
    - **pcmpath query wwpn**
- Linux:

  Additions to supported platforms. See "Software" on page 171 for a complete list of supported platforms.
- Novell NetWare:
  - Addition of support for NetWare 6.5 SP 1.1 operating system that is attached to an ESS.
- Sun Solaris:
  - Addition of support for SAN Volume Controller on a SPARC system running 64-bit Solaris 9.
- Datapath commands:
  - Support for the following new datapath commands:
    - **datapath disable ports** (AIX host systems only)
    - **datapath enable ports** (AIX host systems only)

- **datapath query essmap** (AIX host systems only)
- **datapath query portmap** (AIX host systems only)
- **datapath query wwpn** (AIX, Linux, Windows NT, and Windows 2000 host systems only)
- **datapath remove adapter** (AIX host systems only)
- **datapath remove device path** (AIX host systems only)
- **datapath set qdepth** (AIX host systems only)

• Changes to sddsrv operation.

Discussion in this edition supports:

• IBM TotalStorage Enterprise Storage Server (ESS)
• IBM TotalStorage SAN Volume Controller
• IBM TotalStorage SAN Volume Controller for Cisco MDS 9000
• IBM TotalStorage SAN File System (based on IBM Storage Tank™ technology)

unless the sections are clearly labeled to support one or the other.

Modified information

This edition includes the following modified information:

• The changes to Linux kernel support include:
  – Support for SuSE SLES 7 kernel k_smp-2.4.18-224.
  – Support for the following SLES8 / UnitedLinux 1.0 kernels on i686:
    - k_smp-2.4.21-138
    - k_smp-2.4.21-169
    - k_smp-2.4.21-190
    - k_smp-2.4.21-198
  – Addition of support for the following kernels on pSeries:
    - SLES8 kernel-ppc64-2.4.21-111
    - SLES8 kernel-ppc64-2.4.19-147
    - SLES8 kernel-ppc64-2.4.19-154
  – Red Hat 7.x kernel-2.4.18-24.7.xbigmem
  – Red Hat 7.xkernel-2.4.18-27.7.xbigmem
  – Red Hat Linux Advanced Server 2.1 with:
    - 2.4.9-e.30enterprise
    - 2.4.9-e.34enterprise
    - 2.4.9-e.35enterprise
    - 2.4.9-e.38enterprise
  – Red Hat Linux Advanced Server 2.1 with:
    - 2.4.9-3.30summit
    - 2.4.9-3.34summit
    - 2.4.9-3.35summit
    - 2.4.9-e.38summit

  - Red Hat EL3 with:
    - 2.4.21-9.ELsmp
    - 2.4.21-9.ELhugemem
    - 2.4.21-9.0.1.ELsmp
    - 2.4.21-9.0.1.ELhugemem
  - Red Hat EL3 with:
    - 2.4.21-9.EL
    - 2.4.21-9.0.1.EL
- Modification to existing datapath commands:
  - Addition of suspend and resume options for the **datapath set adapter** command (AIX only)
  - Addition of suspend and resume options for the **datapath open device path** command (AIX only)
- Corrections as necessary.

Related information

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

- The publications that make up the IBM® TotalStorage® Enterprise Storage Server® (ESS) library
- The publications that make up the IBM TotalStorage SAN Volume Controller library
- The publications that make up the IBM TotalStorage SAN Volume Controller for Cisco MDS 9000 library
- The publications that make up the IBM TotalStorage SAN File System library
- Other IBM publications that relate to the ESS
- Non-IBM publications that relate to the ESS

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

The ESS library

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. These publications are also available as PDF files by clicking on the **Documentation link** on the following ESS Web site:

www-1.ibm.com/servers/storage/support/disk/2105.html

See “IBM publications center” on page xxx 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 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 @server 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® and IBM @server 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 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
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® (AS/400®), and @server 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 logical unit numbers. SDD permits applications to run without interruption when path errors occur. It balances the workload across paths, and it transparently integrates with applications. For information about SDD, go to the following Web site: www-1.ibm.com/servers/storage/support/software/sdd.html	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.)

Title	Description	Order Number
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, allowing 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 Application Programming Interface.	GC35-0489
Note: No hardcopy book is produced for this publication. However, a PDF file is available from the following Web site: www-1.ibm.com/servers/storage/support/disk/2105.html		

The SAN Volume Controller Library

The following publications make up the IBM SAN Volume Controller library. These publications are available as Adobe PDFs on a compact disc (CD) that comes with the SAN Volume Controller.

Table 1. Publications in the SAN Volume Controller library

Title	Order number
IBM TotalStorage Virtualization Family: SAN Volume Controller Installation Guide	SC26-7541-00
IBM TotalStorage Virtualization Family: SAN Volume Controller Service Guide	SC26-7542-00
IBM TotalStorage Virtualization Family: SAN Volume Controller Configuration Guide	SC26-7543-00
IBM TotalStorage Virtualization Family: SAN Volume Controller Command-Line Interface User's Guide	SC26-7544-00
IBM TotalStorage Virtualization Family: SAN Volume Controller Planning Guide	GA22-1052-00
IBM TotalStorage Virtualization Family: SAN Volume Controller CIM Agent Developer's Reference	SC26-7545-00
IBM TotalStorage Virtualization Family: SAN Volume Controller Host Systems Attachment Guide	SC26-7563-00

The SAN Volume Controller for Cisco MDS 9000 library

The following publications make up the IBM SAN Volume Controller for Cisco MDS 9000 library. These publications are available as Adobe PDFs on a CD-ROM that comes with the SAN Volume Controller for Cisco MDS 9000.

Table 2. Publications in the SAN Volume Controller for Cisco MDS 9000 library

Title	Order number
IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Installation Guide	SC26-7552-00
IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Service Guide	SC26-7553-00
IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Configuration Guide	SC26-7554-00
IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Command-Line Interface User's Guide	SC26-7555-00
IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Planning Guide	GA22-1055-00
IBM TotalStorage Virtualization Family: SAN Volume Controller Host Systems Attachment Guide	SC26-7563-00

The SAN File System library

Table 3 shows the publications that are available in softcopy form in the SAN File System library.

Table 3. Publications in the SAN File System library

Title	Description	Order Number
IBM TotalStorage SAN File System Release Notes	This document provides any changes that were not available at the time the publications were produced. This document is available only from the technical support Web site: www.ibm.com/storage/support/	
IBM TotalStorage SAN File System License Information	This publication provides multilingual information regarding the software license for IBM TotalStorage SAN File System software.	GC30-9703
IBM TotalStorage SAN File System Administrator's Guide and Reference	This publication introduces the concept of SAN File System and provides instructions for configuring, managing, and monitoring the system using the SAN File System console and administrative command-line interfaces. This book also contains a commands reference for tasks that can be performed at the administrative command-line interface or the command window on the client machines.	GA27-4317
IBM TotalStorage SAN File System Basic Configuration for a Quick Start	This document walks you through basic SAN File System configuration and specific tasks that exercise basic SAN File System functions. It assumes that the physical configuration and software setup have already been completed.	GX27-4058
IBM TotalStorage SAN File System Installation and Configuration Guide	This publication provides detailed procedures to set up and cable the hardware, install and upgrade the SAN File System software, perform the minimum required configuration, and migrate existing data.	GA27-4316
IBM TotalStorage SAN File System Maintenance and Problem Determination Guide	This publication provides instructions for adding and replacing hardware components, upgrading software, monitoring and troubleshooting the system, and resolving hardware and software problems. **Note:** This document is intended only for trained service personnel.	GA27-4318

Table 3. Publications in the SAN File System library (continued)

Title	Description	Order Number
IBM TotalStorage SAN File System Messages Reference	This publication contains message description and resolution information for errors that can occur in theSAN File System software.	GC30-4076
IBM TotalStorage SAN File System Planning Guide	This publication provides detailed procedures to plan the installation and configuration of SAN File System.	GA27-4344
IBM TotalStorage SAN File System System Management API Guide and Reference	This publication contains guide and reference information for using the CIM Proxy application programming interface (API), including common and SAN File System-specific information.	GA27-4315
Note: The softcopy version of this guide and the other related publications are accessibility-enabled for the IBM Home Page Reader. The softcopy publications support the SAN File System system and they are available on the *IBM TotalStorage SAN File System CD* that came with your appliance and at www.ibm.com/storage/support/		

Other IBM publications

Other IBM publications contain additional information that is related to the ESS. The following list is divided into categories to help you find publications that are related to specific topics. Some of the publications are listed under more than one category. See “IBM publications center” on page xxx for information about ordering these and other IBM publications.

Table 4. Other IBM publications

Title	Description	Order Number
Data-copy services		
DFSMS/MVS® Version 1: Advanced Copy Services	This publication helps you understand and use IBM Advanced Copy Services functions on an S/390 or zSeries. It describes two dynamic copy functions and several point-in-time copy functions. These functions provide backup and recovery of data if a disaster occurs to your data center. The dynamic copy functions are peer-to-peer remote copy and extended remote copy. Collectively, these functions are known as remote copy. FlashCopy and concurrent copy are the point-in-time copy functions.	SC35-0355
DFSMS/MVS V1: Remote Copy Guide and Reference	This publication provides guidelines for using remote copy functions with S/390 and zSeries hosts.	SC35-0169
IBM Enterprise Storage Server	This publication, from the IBM International Technical Support Organization, introduces the ESS and provides an understanding of its benefits. It also describes in detail the architecture, hardware, and functions, including the advanced copy functions, of the ESS.	SG24-5465
Implementing Copy Services 0n S/390	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an ESS that is attached to an S/390 or zSeries host system. Copy Services functions include peer-to-peer remote copy, extended remote copy, FlashCopy, and concurrent copy. This publication describes the functions, prerequisites, and corequisites and describes how to implement each function into your environment.	SG24-5680

Table 4. Other IBM publications (continued)

Title	Description	Order Number
IBM TotalStorage ESS Implementing Copy Services in an Open Environment	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure ESS Copy Services on UNIX, Windows NT, Windows 2000, Sun Solaris, HP-UX, Tru64, Open VMS, and iSeries host systems. The Copy Services functions that are described include peer-to-peer remote copy and FlashCopy. This publication describes the functions and shows you how to implement them into your environment. It also shows you how to implement these functions in a high-availability cluster multiprocessing environment.	SG24-5757
Fibre channel		
Fibre Channel Connection (FICON) I/O Interface: Physical Layer	This publication provides information about the fibre-channel I/O interface.	SA24-7172
Fibre-channel Subsystem Installation Guide	This publication tells you how to attach the IBM @server xSeries® 430 and NUMA-Q® host systems with fibre-channel adapters to the ESS. Contact your sales representative to obtain this publication.	No order number
Fibre Transport Services (FTS): Physical and Configuration Planning Guide	This publication provides information about fiber-optic and IBM ESCON® trunking systems.	GA22-7234
IBM SAN Fibre Channel Managed Hub: 3534 Service Guide	This guide explains how to convert the IBM SAN Fibre Channel Managed Hub from a Fibre-Channel Arbitrated Loop (FC-AL) configuration to a fabric-capable switched environment.	SY27–7616
IBM SAN Fibre Channel Managed Hub: 3534 User's Guide	This guide provides an overview of the product and discusses available features and upgrades. It also tells you how to install the hub, how to manage and monitor it using zoning, and how to manage it remotely. It also tells you how to use the IBM 3534 SAN Fibre Channel Managed Hub TotalStorage ESS Specialist.	GC26-7391
IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08. To get a copy of this guide, go to the following Web site: www.ibm.com/storage/ibmsan/products/2109/san_switch_solu.html	SC26-7350
IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide	This guide describes the IBM SAN Fibre Channel Switch and the IBM TotalStorage ESS Specialist. It provides information about the commands and how to manage the switch with Telnet and the Simple Network Management Protocol. To get a copy of this guide, go to the following Web site: www.ibm.com/storage/ibmsan/products/2109/san_switch_solu.html	SC26-7349
IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide	This publication describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S16. It is intended for trained service representatives and service providers. To get a copy of this guide, go to the following Web site: www.ibm.com/storage/ibmsan/products/2109/san_switch_solu.html	SC26-7352

Table 4. Other IBM publications (continued)

Title	Description	Order Number
IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide	This guide introduces the IBM SAN Fibre Channel Switch 2109 Model S16 and tells you how to manage and monitor the switch using zoning and how to manage the switch remotely. To get a copy of this guide, go to the following Web site: www.ibm.com/storage/ibmsan/products/2109/san_switch_solu.html	SC26-7351
Implementing Fibre Channel Attachment on the ESS	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the ESS. It provides you with a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment. This book also describes the steps required to migrate to direct fibre-channel attachment from native SCSI adapters and from fibre-channel attachment through the SAN Data Gateway.	SG24-6113
Open-systems hosts		
ESA/390: ESCON I/O Interface	This publication provides a description of the physical and logical ESA/390 I/O interface and the protocols that govern information transfer over that interface. It is intended for designers of programs and equipment associated with the ESCON I/O interface and for service personnel who maintain that equipment. However, anyone concerned with the functional details of the ESCON I/O interface can find it useful.	SA22-7202
ESS Solutions for Open Systems Storage: Compaq AlphaServer, HP, and Sun	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure the ESS when you attach Compaq AlphaServer (running Tru64 UNIX®), HP, and Sun hosts. This book does not cover Compaq AlphaServer that is running the Open VMS operating system. This book also focuses on the settings that are required to give optimal performance and on the settings for device driver levels. This book is for the experienced UNIX professional who has a broad understanding of storage concepts.	SG24-6119
Fibre-channel Subsystem Installation Guide	This publication tells you how to attach the IBM xSeries 430 and NUMA-Q host systems with fibre-channel adapters to the ESS. Contact your sales representative to obtain this publication.	No order number
IBM TotalStorage ESS Implementing Copy Services in an Open Environment	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure ESS Copy Services on UNIX, Windows NT, or Windows 2000 host systems. The Copy Services functions that are described include peer-to-peer remote copy and FlashCopy. This publication describes the functions and shows you how to implement them into your environment. It also shows you how to implement these functions in a high-availability cluster multiprocessing environment.	SG24-5757
Implementing Fibre Channel Attachment on the ESS	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the ESS. It gives you a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment. This book also describes the steps that are required to migrate to direct fibre-channel attachment from native SCSI adapters and from fibre-channel attachment through the SAN Data Gateway.	SG24-6113
S/390 and zSeries hosts		

Table 4. Other IBM publications (continued)

Title	Description	Order Number
Device Support Facilities: User's Guide and Reference	This publication describes the IBM Device Support Facilities (ICKDSF) product that is used with IBM direct access storage device (DASD) subsystems. ICKDSF is a program that you can use to perform functions that are needed for the installation, the use, and the maintenance of IBM DASD. You can also use it to perform service functions, error detection, and media maintenance.	GC35-0033
DFSMS/MVS Version 1: Advanced Copy Services	This publication helps you understand and use IBM Advanced Copy Services functions on an S/390 or zSeries host. It describes two dynamic copy functions and several point-in-time copy functions. These functions provide backup and recovery of data if a disaster occurs to your data center. The dynamic copy functions are peer-to-peer remote copy and extended remote copy. Collectively, these functions are known as remote copy. FlashCopy and concurrent copy are the point-in-time copy functions.	SC35-0355
DFSMS/MVS V1: Remote Copy Guide and Reference	This publication provides guidelines for using remote copy functions with S/390 and zSeries hosts.	SC35-0169
ESA/390: ESCON I/O Interface	This publication provides a description of the physical and logical ESA/390 I/O interface and the protocols that govern information transfer over that interface. It is intended for designers of programs and equipment associated with the ESCON I/O interface and for service personnel who maintain that equipment. However, anyone concerned with the functional details of the ESCON I/O interface will find it useful.	SA22-7202
Fibre Transport Services (FTS): Physical and Configuration Planning Guide	This publication provides information about fiber-optic and ESCON-trunking systems.	GA22-7234
Implementing ESS Copy Services on S/390	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an ESS that is attached to an S/390 or zSeries host system. Copy Services functions include Peer-to-Peer Remote Copy, Extended Remote Copy, FlashCopy, and Concurrent Copy. This publication describes the functions, prerequisites, and corequisites and describes how to implement each function into your environment.	SG24-5680
ES/9000, ES/3090: IOCP User Guide Volume A04	This publication describes the Input/Output Configuration Program that supports the Enterprise Systems Connection (ESCON) architecture. It describes how to define, install, and configure the channels or channel paths, control units, and I/O devices on the IBM ES/9000® processors and the IBM ES/3090™ Processor Complex.	GC38-0097
IOCP User's Guide, IBM @server zSeries	This publication describes the Input/Output Configuration Program that supports the zSeries server.	SB10-7029
S/390: Input/Output Configuration Program User's Guide and ESCON Channel-to-Channel Reference	This publication describes the Input/Output Configuration Program that supports ESCON architecture and the ESCON multiple image facility.	GC38-0401
IBM z/OS Hardware Configuration Definition User's Guide	This guide provides conceptual and procedural information to help you use the z/OS® Hardware Configuration Definition (HCD) application. It also explains: • How to migrate existing IOCP/MVSCP definitions • How to use HCD to dynamically activate a new configuration • How to resolve problems in conjunction with MVS/ESA™ HCD	SC33-7988

Table 4. Other IBM publications (continued)

Title	Description	Order Number
OS/390: Hardware Configuration Definition User's Guide	This guide provides detailed information about the input/output definition file and about how to configure parallel access volumes. This guide discusses how to use Hardware Configuration Definition for both OS/390® and z/OS V1R1.	SC28-1848
OS/390 V2R10.0: MVS System Messages Volume 1 (ABA - ASA)	This publication lists OS/390 MVS™ system messages ABA to ASA.	GC28-1784
Using IBM 3390 Direct Access Storage in a VM Environment	This publication provides device-specific information for the various models of the 3390 and describes methods you can use to manage storage efficiently using the VM operating system. It provides guidance on managing system performance, availability, and space through effective use of the direct access storage subsystem.	GG26-4575
Using IBM 3390 Direct Access Storage in a VSE Environment	This publication helps you use the 3390 in a VSE environment. It includes planning information for adding new 3390 units and instructions for installing devices, migrating data, and performing ongoing storage management activities.	GC26-4576
Using IBM 3390 Direct Access Storage in an MVS Environment	This publication helps you use the 3390 in an MVS environment. It includes device-specific information for the various models of the 3390 and illustrates techniques for more efficient storage management. It also offers guidance on managing system performance, availability, and space utilization through effective use of the direct access storage subsystem.	GC26-4574
z/Architecture Principles of Operation	This publication provides a detailed definition of the IBM z/Architecture™. It is written as a reference for use primarily by assembler language programmers and describes each function at the level of detail needed to prepare an assembler language program that relies on a particular function. However, anyone concerned with the functional details of z/Architecture will find this publication useful.	SA22-7832
SAN		
IBM OS/390 Hardware Configuration Definition User's Guide	This guide explains how to use the Hardware Configuration Data application to perform the following tasks: • Define new hardware configurations • View and modify existing hardware configurations • Activate configurations • Query supported hardware • Maintain input/output definition files (IODFs) • Compare two IODFs or compare an IODF with an actual configuration • Print reports of configurations • Create graphical reports of a configuration • Migrate existing configuration data	SC28-1848
IBM SAN Fibre Channel Managed Hub: 3534 Service Guide	This guide explains how to convert the IBM SAN Fibre Channel Managed Hub from a Fibre-Channel Arbitrated Loop (FC-AL) configuration to a fabric-capable switched environment.	SY27–7616
IBM SAN Fibre Channel Managed Hub: 3534 User's Guide	This guide provides an overview of the product and discussed the features and upgrades available. It also tells you how to install the hub, how to manage and monitor it using zoning, and how to manage it remotely. It also tells you how to use the IBM 3534 SAN Fibre Channel Managed Hub TotalStorage ESS Specialist.	GC26-7391

Table 4. Other IBM publications (continued)

Title	Description	Order Number
IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08. To get a copy of this guide, go to the following Web site: www.ibm.com/storage/ibmsan/products/2109/san_switch_solu.html	SC26-7350
IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide	This guide describes the IBM SAN Fibre Channel Switch and the IBM TotalStorage ESS Specialist. It provides information about the commands and how to manage the switch with Telnet and the Simple Network Management Protocol (SNMP). To get a copy of this guide, go to the following Web site: www.ibm.com/storage/ibmsan/products/2109/san_switch_solu.html	SC26-7349
IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide	This publication describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S16. It is intended for trained service representatives and service providers. To get a copy of this guide, go to the following Web site: www.ibm.com/storage/ibmsan/products/2109/san_switch_solu.html	SC26-7352
IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide	This guide introduces the IBM SAN Fibre Channel Switch 2109 Model S16 and tells you how to manage and monitor the switch using zoning and how to manage the switch remotely. To get a copy of this guide, go to the following Web site: www.ibm.com/storage/ibmsan/products/2109/san_switch_solu.html	SC26-7351
Implementing Fibre Channel Attachment on the ESS	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the ESS. It provides you with a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment. This book also describes the steps required to migrate to direct fibre-channel attachment from native SCSI adapters and from fibre-channel attachment through the SAN Data Gateway.	SG24-6113
Seascape family		
IBM Enterprise Storage Server	This publication, from the IBM International Technical Support Organization, introduces the ESS and provides an understanding of its benefits. It also describes in detail the architecture, hardware, and functions, including the advanced copy functions, of the ESS.	SG24-5465
IBM Enterprise Storage Server Performance Monitoring and Tuning Guide	This guide, from the IBM International Technical Support Organization, provides guidance on the best way to configure, monitor, and manage your ESS to ensure optimum performance.	SG24-5656
IBM Versatile Storage Server: Introduction and Planning Guide	This publication introduces the IBM Versatile Storage Server™ and lists the features you can order. It also provides planning information for both 2105 Models B09 and 100.	GC26-7223
Implementing the IBM Enterprise Storage Server in Your Environment	This publication, from the IBM International Technical Support Organization, can help you install, tailor, and configure the ESS in your environment.	SG24-5420
Storage management		

Table 4. Other IBM publications (continued)

Title	Description	Order Number
Device Support Facilities: User's Guide and Reference	This publication describes the IBM Device Support Facilities (ICKDSF) product used with IBM direct access storage device (DASD) subsystems. ICKDSF is a program that you can use to perform functions that are needed for the installation, the use, and the maintenance of IBM DASD. You can also use it to perform service functions, error detection, and media maintenance.	GC35-0033
IBM TotalStorage Solutions Handbook	This handbook, from the IBM International Technical Support Organization, helps you understand what makes up enterprise storage management. The concepts include the key technologies that you must know and the IBM subsystems, software, and solutions that are available today. It also provides guidelines for implementing various enterprise storage administration tasks so that you can establish your own enterprise storage management environment.	SG24-5250
IBM TotalStorage Expert: Hands-On Usage Guide	This guide, from the IBM International Technical Support Organization, helps you install, tailor, configure, and use TotalStorage ESS Expert.	SG24-6102
IBM TotalStorage Expert Installation Guide	This guide helps you install the IBM TotalStorage Expert (formerly the IBM StorWatch Expert) program. The IBM TotalStorage Expert provides asset, capacity, and performance management information for disk and tape storage systems.	GC26-7436
Using IBM 3390 Direct Access Storage in a VM Environment	This publication provides device-specific information for the various models of the 3390 and describes methods that you can use to manage storage efficiently using the VM operating system. It provides guidance for managing system performance, availability, and space through effective use of the direct access storage subsystem.	GG26-4575
Using IBM 3390 Direct Access Storage in a VSE Environment	This publication helps you use the 3390 in a VSE environment. It includes planning information for adding new 3390 units and instructions for installing devices, migrating data, and performing ongoing storage management activities.	GC26-4576
Using IBM 3390 Direct Access Storage in an MVS Environment	This publication helps you use the 3390 in an MVS environment. It includes device-specific information for the various models of the 3390 and illustrates techniques for more efficient storage management. It also offers guidance for managing system performance, availability, and space use through effective use of the direct access storage subsystem.	GC26-4574

Ordering IBM publications

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

IBM publications center

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

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

www.ibm.com/shop/publications/order/

Publications notification system

The IBM publications center Web site offers you a notification system for IBM publications. Register and you can create your own profile of publications that interest you. The publications notification system sends you a daily e-mail that contains information about new or revised publications that are based on your profile.

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

www.ibm.com/shop/publications/order/

Non-IBM publications

Other related publications are not available through IBM ordering systems. To order them, contact the sales representative at the branch office in your locality.

Title	Description
Quick Start Guide: An Example with Network File System (NFS)	This guide tells you how to configure the Veritas Cluster Server. See also the companion document, *Veritas Cluster Server User's Guide*.
Veritas Cluster Server Installation Guide	This guide tells you how to install the Veritas Cluster Server. See also the companion document, *Veritas Cluster Server Release Notes*.
Veritas Cluster Server Release Notes	These release notes tell you how to install the Veritas Cluster Server. See also the companion document, *Veritas Cluster Server Installation Guide*.
Veritas Cluster Server User's Guide	This guide tells you how to configure the Veritas Cluster Server. See also the companion document, *Quick Start Guide: An Example with Network File System (NFS)*.
Veritas Volume Manager Hardware Notes	These hardware notes tell you how to implement multiple paths dynamically.
Veritas Volume Manager Installation Guide	This guide tells you how to install VxVM.
Veritas Volume Manager Storage Administrators Guide	This guide tells you how to administer and configure the disk volume groups.

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
  - Internet: starpubs@us.ibm.com
  - IBMLink™ from U.S.A: STARPUBS at SJEVM5
  - IBMLink from Canada: STARPUBS at TORIBM
  - IBM Mail Exchange: USIB3WD at IBMMAIL

  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

  Fill out the Readers' Comments form (RCF) at the back of this book. Return it by mail or fax (1-800-426-6209), or give it to an IBM representative. If the RCF has been removed, you can address your comments to:

  International Business Machines Corporation
  Information Development
  Department GZW
  9000 South Rita Road
  Tucson, AZ 85744-0001
  U.S.A.

Chapter 1. Overview of SDD

This chapter introduces the IBM Subsystem Device Driver (SDD) and provides an overview of the SDD functions.

The SDD architecture

SDD is a pseudo device driver designed to support the multipath configuration environments in IBM products. 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 internal code

Table 5 shows the position of SDD in the protocol stack. I/O operations that are sent to SDD proceed to the host disk driver after path selection. When an active path experiences a failure (such as a cable or controller failure), SDD switches to another path dynamically.

Table 5. SDD in the protocol stack

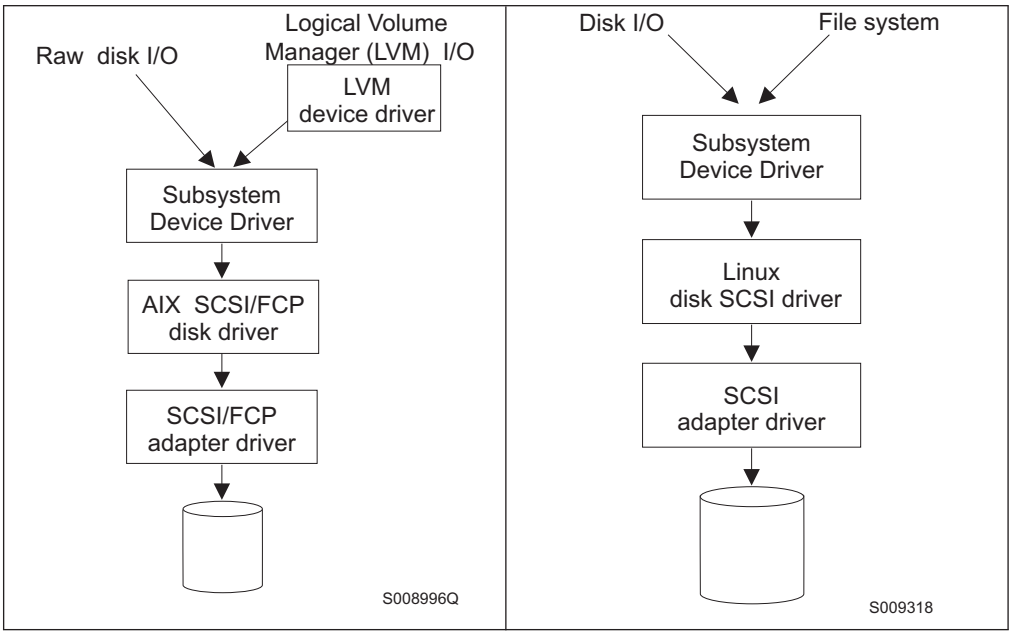

© Copyright IBM Corp. 1999, 2004

Table 5. SDD in the protocol stack (continued)

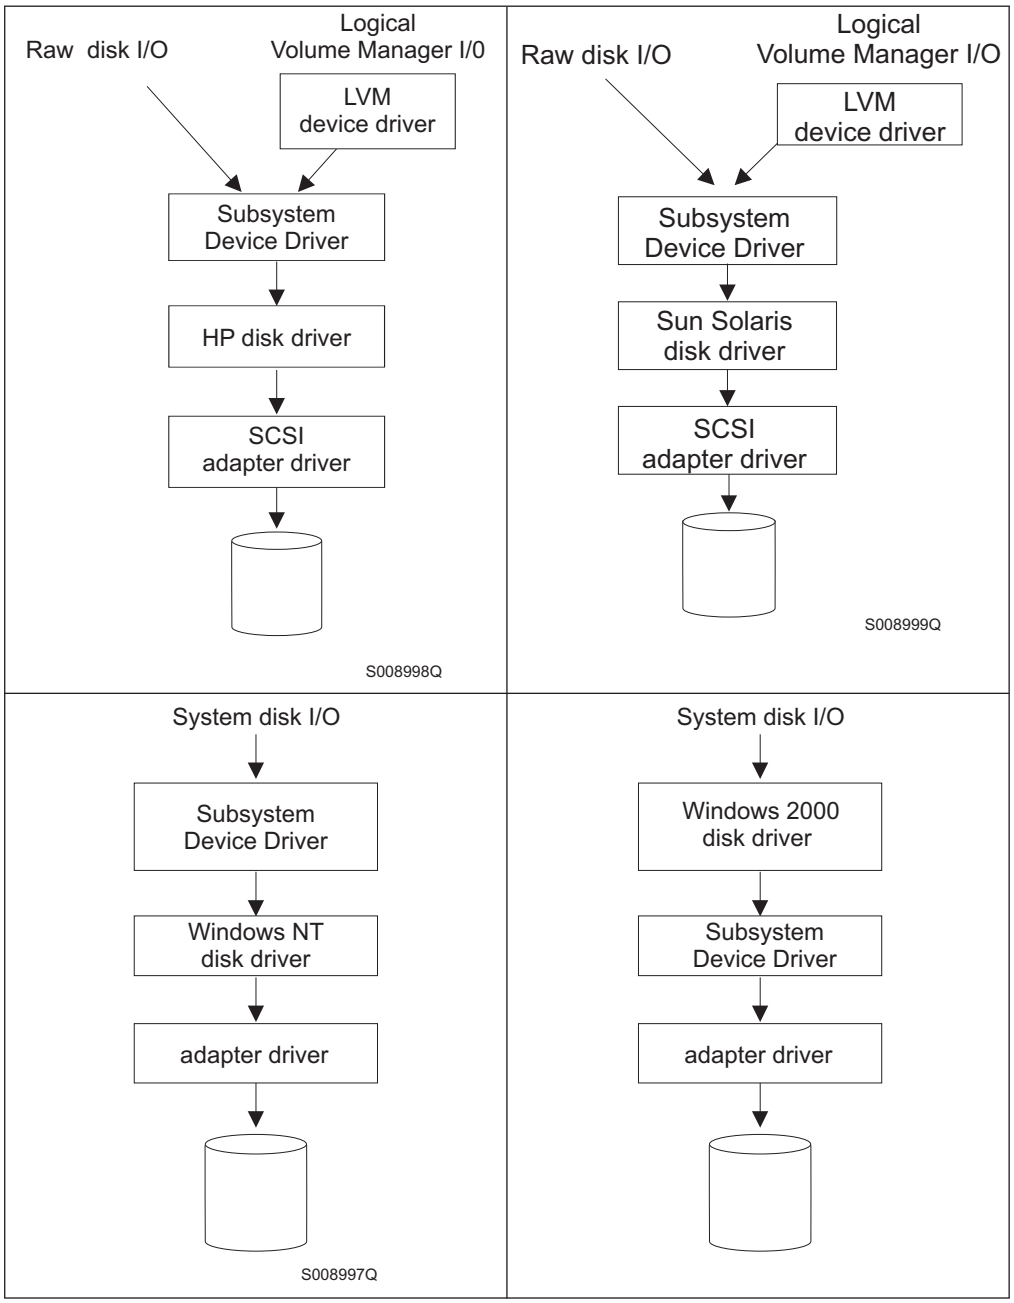

Table 5. SDD in the protocol stack (continued)

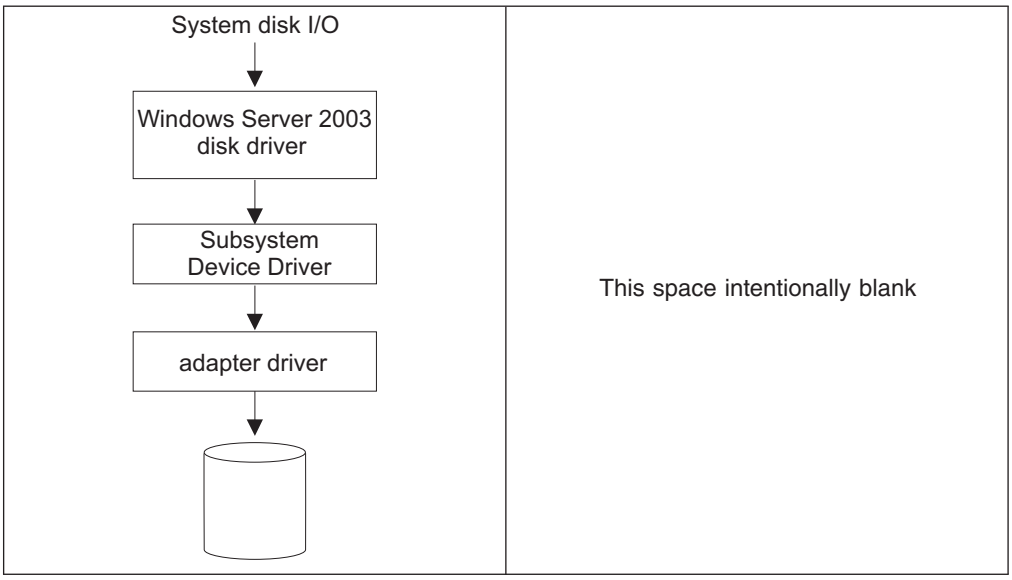

Each SDD device represents a unique physical device on the storage server. Each physical device is presented to the operating system as an operating system disk device. There can be up to 32 operating system disk devices that represent up to 32 different paths to the same physical device.

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

Note: Commands getting the information directly from the system file where the device driver is not involved for the operation are not supported for the SDD devices.

Enhanced data availability

Figure 1 on page 4 shows a host system that is attached through small computer system interface (SCSI) or fibre-channel adapters to an ESS that has internal component redundancy and multipath configuration. SDD, residing in the host system, uses this multipath configuration to enhance data availability. That is, when there is a path failure, 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 ESS from disrupting data access.

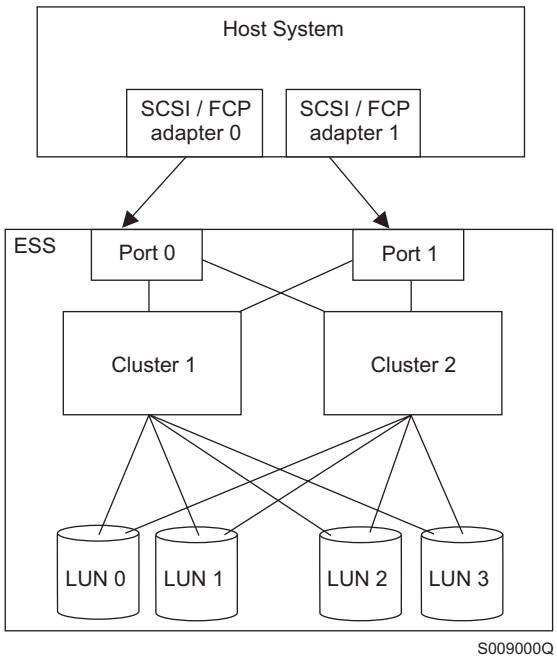

Figure 1. Multipath connections between a host system and the disk storage in an ESS

Figure 2 shows a host system that is attached through fibre-channel adapters to a SAN Volume Controller that has internal component redundancy and multipath configuration. SDD, residing in the host system, uses this multipath configuration to enhance data availability. That is, when there is a path failure, 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.

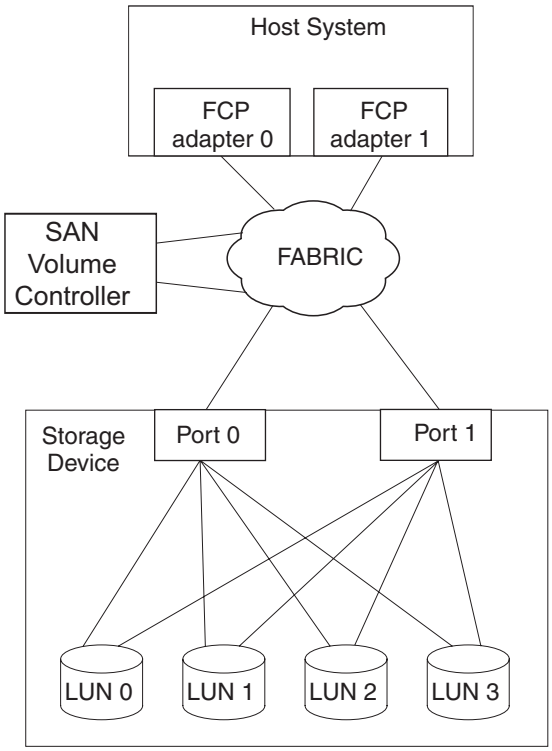

Figure 2. Multipath connections between a host system and the disk storage with a SAN Volume Controller

Figure 3 shows a host system that is attached through fibre-channel adapters to a SAN Volume Controller for Cisco MDS 9000 that has internal component redundancy and multipath configuration. SDD, residing in the host system, uses this multipath configuration to enhance data availability. That is, when there is a path failure, 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 for Cisco MDS 9000 from disrupting data access.

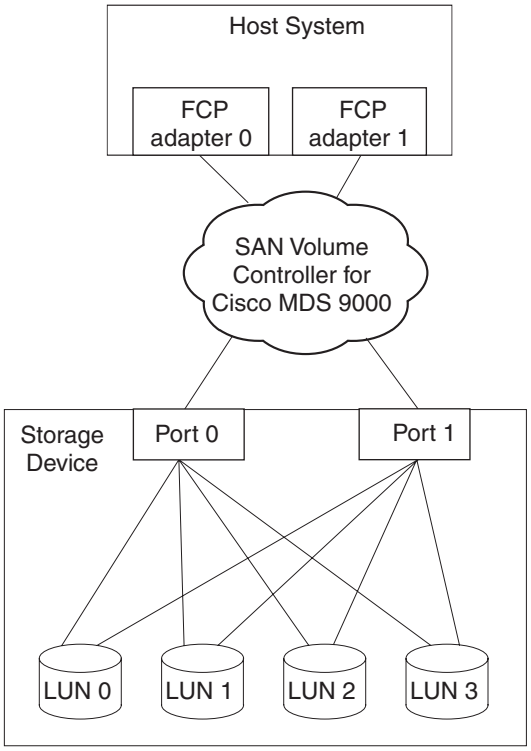

Figure 3. Multipath connections between a host system and the disk storage with a SAN Volume Controller for Cisco MDS 9000

Note: The SAN Volume Controller and the SAN Volume Controller for Cisco MDS 9000 do not support parallel SCSI attachment.

Dynamic I/O load balancing

By distributing the I/O workload over multiple active paths, SDD provides dynamic load balancing and eliminates data-flow bottlenecks. In the event of failure in one data path, 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 system minimizes any disruptions in I/O operations and recovers I/O operations from a failing data path. SDD provides path-failover protection by using the following process:

- Detect a path failure.
- Notify the host system of the path failure.
- Select and use an alternate data path.

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 SDD's automatic path-failover feature, some reported errors are actually recovered from an alternative path.

Concurrent download of ESS licensed internal code

With SDD multipath mode (configured with at least two paths per SDD device), you can concurrently download and install the licensed internal code (LIC) while applications continue to run. For certain ESS LIC, the ESS bay will be quiesced and resumed. Its adapters might not respond to host I/O requests for approximately 30 minutes.

Note: SDD does not support single-path mode during the concurrent download and installation of LIC. Also, SDD does not support single-path mode during any ESS concurrent maintenance that impacts the path attachment, such as an ESS host-bay-adapter replacement.

For information about performing concurrent download of LIC for the ESS, refer to the Microcode Installation Instructions for your specific type and model.

Concurrent download of SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 licensed internal code

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

At least one path must be configured through each node of a SAN Volume Controller or a SAN Volume Controller for Cisco MDS 9000 I/O group. That is, if only two paths exist, then they must go to separate nodes for each I/O group. However, at least 2 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 should not see any I/O failures.

For information about performing concurrent download of LIC for the SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000, refer to the Configuration Guide for your specific type and model.

Chapter 2. Using SDD on an AIX host system

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

Starting from level 1.4.0.5, SDD supports the coexistence of ESS devices and SAN Volume Controller devices.

SDD 1.5.0.0 or later supports the coexistence of ESS devices and SAN Volume Controller for Cisco MDS 9000.

SAN File System might require a specific version of SDD for multipathing support to ESS and SAN Volume Controller devices. Refer to SAN File System documentation shown in Table 3 on page xxiii for the latest information about the version of SDD that is required by SAN File System.

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

www-1.ibm.com/servers/storage/support/software/sdd.html

Verifying the hardware and software requirements

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

Hardware

- One or more supported storage devices.
- A switch if using a SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 (no direct attachment allowed for SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000)
- Host system
- SCSI adapters and cables (ESS only)
- Fibre-channel adapters and cables

Software

- AIX operating system
- 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.ibm2145.rte package for SAN Volume Controller devices
- devices.fcp.disk.ibm2062.rte package for SAN Volume Controller for Cisco MDS 9000 devices

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 6 on page 10 shows the package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0 (or later).

© Copyright IBM Corp. 1999, 2004

Table 6. 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	N/A	Obsolete. This package has been merged with devices.sdd.43.rte.
ibmSdd_433.rte	devices.sdd.43.rte	N/A
ibmSdd_510.rte	N/A	Obsolete. This package has been merged with devices.sdd.51.rte.
ibmSdd_510nchacmp.rte	devices.sdd.51.rte	N/A
N/A	devices.sdd.52.rte	New package for AIX 5.2.0 (or later).

Notes:

1. SDD 1.4.0.0 (or later) no longer releases separate packages for concurrent and nonconcurrent High Availability Cluster Multi-Processing (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 52 for more information.

Unsupported environments

SDD does not support:

- A host system with both a SCSI and fibre-channel connection to a shared ESS logical unit number (LUN)
- A system restart from an SDD pseudo device
- Placing system primary paging devices (for example, /dev/hd6) on an SDD pseudo device
- Any application that depends on a SCSI-2 reserve and release device on AIX
- Single-path mode during concurrent download and installation of LIC nor during any ESS concurrent maintenance that impacts the path attachment, such as an ESS host-bay-adapter replacement
- Multipathing to an ESS boot device (rootvg with ESS device)
- Dump devices–you should not configure an SDD vpath device as a secondary dump device
- More than 600 vpaths if the host system is running AIX 4.3.2, AIX 4.3.3 or AIX 5.1.0

Host system requirements

To successfully install SDD for ESS and for SAN Volume Controller, you must have AIX 4.3, AIX 5.1, or AIX 5.2 installed on your host system.

To successfully install SDD for SAN Volume Controller for Cisco MDS 9000, you must have AIX 5.1 or AIX 5.2 installed on your host system.

SAN File System might require a specific version of SDD for multipathing support. Refer to the SAN File System documentation shown in “The SAN File System library” on page xxiii for the latest information about the version of SDD that is required by SAN File System.

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

www-1.ibm.com/servers/eserver/support/pseries/fixes/

ESS requirements

To successfully install SDD, ensure that the ESS devices are configured as either an:

– IBM 2105xxx (SCSI-attached device)

– IBM FC 2105xxx (fibre-channel-attached device)

where *xxx* represents the ESS model number.

SAN Volume Controller requirements

To successfully install SDD, ensure that the SAN Volume Controller devices are configured as:

SAN Volume Controller Device (fibre-channel-attached device)

SAN Volume Controller for Cisco MDS 9000 requirements

To successfully install SDD, ensure that the SAN Volume Controller for Cisco MDS 9000 devices are configured as:

SVCCISCO Device (fibre-channel-attached device)

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.

For information about the SCSI adapters that can attach to your AIX host system, go to the following Web site:

www-1.ibm.com/servers/eserver/support/

Fibre requirements

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

www-1.ibm.com/servers/eserver/support/

Notes:

1. If your host has only one fibre-channel adapter, it requires you to connect through a switch to multiple ESS ports. You should 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 whether or not the host has only one fibre-channel adapter. Refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Installation Guide*.

3. The SAN Volume Controller for Cisco MDS 9000 ***always*** requires that the host be connected directly to it. The SAN Volume Controller for Cisco MDS 9000 cannot be cascaded through another storage device. For information about the SAN Volume Controller for Cisco MDS 9000, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Installation Guide*.

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

www-1.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 pSeries 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 an ESS port.
- A fiber-optic cable connects each SAN Volume Controller fibre-channel adapter to a switch. The switch must also be configured correctly. Refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller 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 SDD installation

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

- Configuring the ESS
- Configuring the SAN Volume Controller
- Configuring the SAN Volume Controller for Cisco MDS 9000
- Installing the AIX fibre-channel device drivers
- Configuring fibre-channel-attached devices
- Verifying the adapter firmware level
- Determining if the sddServer for Expert is installed
- Determining the installation package
- Determining the installation type

Note:

- SDD 1.3.3.9 or later supports manual exclusion of ESS devices from the SDD configuration. SDD 1.4.0.0 or later supports manual exclusion of SAN Volume Controller devices from the SDD configuration. SDD 1.5.0.0 or later supports manual exclusion of SAN Volume Controller for Cisco MDS 9000 devices from the SDD configuration.
- If you want to manually exclude supported devices (hdisks) from the SDD configuration, you must do so before configuring SDD devices.
- The **querysn** command can be used to exclude any supported devices (hdisks) from the SDD configuration. It reads the unique serial number of a device (hdisk) and saves the serial number in an exclude file. For

detailed information about the **querysn** command, see “Manual exclusion of devices from the SDD configuration” on page 18.

Configuring the ESS

Before you install SDD, you must configure:

- The ESS to your host system and the required fibre-channel or SCSI adapters that are attached.
- The ESS 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 failover features. **With a single path, failover protection is not provided.**

For more information about how to configure your IBM Enterprise Storage Server, refer to the *IBM TotalStorage Enterprise Storage Server: Introduction and Planning Guide*.

Note: Ensure that the ibm2105.rte installation package is installed.

Configuring the SAN Volume Controller

Before you install SDD, you must configure:

- The SAN Volume Controller to your host system and the required fibre-channel adapters that are attached.
- The SAN Volume Controller for 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.**

For more information about how to configure your IBM SAN Volume Controller, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Configuration Guide*.

Note: Ensure that the devices.fcp.disk.ibm2145.rte installation package is installed before configuring the SAN Volume Controller.

Configuring the SAN Volume Controller for Cisco MDS 9000

Before you install SDD, you must configure:

- The SAN Volume Controller for Cisco MDS 9000 to your host system and the required fibre-channel adapters that are attached.
- The SAN Volume Controller for Cisco MDS 9000 for 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.**

For more information about how to configure your IBM SAN Volume Controller for Cisco MDS 9000, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Configuration Guide*.

Note: Ensure that the devices.fcp.disk.ibm2062.rte installation package is installed before configuring the SAN Volume Controller for Cisco MDS 9000.

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 Web site:

www-1.ibm.com/servers/storage/support/

Perform 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. 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 pSeries with feature code 6228.

   devices.pci.df1000f7
   The adapter device driver for RS/6000 or pSeries 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 pSeries 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." 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

Perform 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

Perform 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**.

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: If operating in a switched environment, the **cfgmgr** command must be executed once for each host adapter each time a device is added.

  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 should be in the *Available* state. If the FCP devices are configured correctly, go to "Verifying the adapter firmware level" on page 16 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 should be in the *Available* state. If the FCP devices are configured correctly, go to "Verifying the adapter firmware level" 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 -dl  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   -dl  fcs0  -R
rmdev   -dl  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 Web site:

http://techsupport.services.ibm.com/server/mdownload/download.html

Tip:

- The current firmware level for the FC 6227 adapter is 3.30X1
- The current firmware level for the FC 6228 adapter is 3.91A1
- The current firmware level for the FC 6239 adapter is 1.00X5

Perform 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).............S2F3.22A1
   ```

   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, refer to the description for each firmware at

   http://techsupport.services.ibm.com/server/mdownload/download.html

Determining if the sddServer for Expert is installed

If you previously installed the *standalone* version of the sddServer for IBM TotalStorage Expert V2R1 (ESS Expert) on your AIX host system, you must remove this standalone 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 standalone version of sddServer (for ESS Expert).

To determine if the standalone 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 instructions on how to remove the standalone version of sddServer (for ESS Expert) from your AIX host system, see the IBM(R) SUBSYSTEM DEVICE DRIVER SERVER 1.0.0.0 (sddsrv) README for IBM TotalStorage Expert V2R1 at the following Web site:

www-1.ibm.com/servers/storage/support/software/swexpert.html

For more information about the SDD server daemon, go to “SDD server daemon” on page 48

Planning for SDD installation on a pSeries 690 server LPAR

As a standard feature, the IBM @server pSeries 690 server supports static logical partitioning (LPARs). The partitions on a pSeries 690 server have their own instances of operating systems, and resources such as processors, dedicated memory and I/O adapters, and they do not share the same hardware resources.

SDD provides the same functions on one of the partitions or LPARs of a pSeries 690 server as it does on a standalone server.

Before you install SDD on one of the partitions or LPARs of a pSeries 690 server, you need to determine the installation package that is appropriate for your environment. See Table 7 to determine the correct installation package.

Determining the installation package

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

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

Table 7. SDD 1.4.0.0 (or later) installation packages for 32-bit and 64-bit applications on AIX 4.3.3 or later

SDD installation package names	AIX kernel level	AIX kernel mode	Application mode	SDD interface
devices.sdd.43.rte	AIX 4.3.2[1], AIX 4.3.3 [2]	32-bit	32-bit, 64-bit	LVM, raw device
devices.sdd.51.rte	AIX 5.1.0 [2]	32-bit, 64-bit	32-bit, 64-bit	LVM, raw device
devices.sdd.52.rte	AIX 5.2.0 [2]	32-bit, 64-bit	32-bit, 64-bit	LVM, raw device

[1] SAN Volume Controller devices are not supported on AIX 4.3.2. SAN Volume Controller for Cisco MDS 9000 devices are not supported on AIX 4.3.2.

[2] SAN File System is not supported on AIX 4.3.3. Refer to the SAN File System documentation shown in "The SAN File System library" on page xxiii for the latest information about the version of SDD that is required by SAN File System.

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

SDD 1.4.0.0 (or later) supports AIX 5.1.0 and AIX 5.2.0 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 or AIX 5.2.0 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 5.1.0 host system runs in 32-bit mode. If the /unix links to /usr/lib/boot/unix_64, your AIX 5.1.0 or AIX 5.2.0 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.

Manual exclusion of devices from the SDD configuration

With the maintenance level of the AIX 4.3.3, AIX 5.1.0, and AIX 5.2.0 operating systems, AIX supports fibre-channel boot capability for selected pSeries 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, you should not configure that device with multipath configuration. Refer to the *IBM TotalStorage Host System Attachment Guide* for the supported storage device for additional information.

The SDD driver will automatically exclude 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, you should use the **querysn** command to manually exclude all boot devices that belong to multiple non-active rootvg volume groups on the server or ESS devices that are going to be selected as physical volumes of a rootvg.

SDD 1.3.3.9 (or later) allows you to manually exclude devices from the SDD configuration. The **querysn** command reads the unique serial number of an device (hdisk) and saves the serial number in an exclude file, /etc/vpexclude. During the SDD configuration, SDD configure methods read all the serial numbers in this exclude file and exclude these devices from the SDD configuration. See “querysn” on page 87 for the syntax of the **querysn** command. The maximum number of devices that can be excluded is 100.

The exclude file, /etc/vpexclude, holds the serial numbers of all devices (hdisks) to be excluded from the SDD configuration in the system. If this exclude file exists, the **querysn** command will add the excluded serial number to that file. If no exclude file exists, the **querysn** command will create one. There is no user interface to this file.

You can also exclude any SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 devices from the SDD configuration with the **querysn** command.

Notes:

1. You should not use the **querysn** command on the same logical device multiple times.
2. Fibre-channel boot capability is available for ESS 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, then you need to perform one of the following actions:
   a. Unmount (umount) 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. Use a text editor such as vi to open the ’/etc/vpexclude’ file and delete the line containing the device name from the file.
3. Execute **cfallvpath** configure methods to configure these new devices.
4. Execute **lsvpcfg** to verify that these devices are configured as SDD vpath devices.

Installation of major files on your AIX host system

The installation package installs a number of major files on your AIX system. Table 8 lists the major files that are part of the SDD installation package.

Table 8. 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 devices while they are in *Available* state.
cfgvpath	Configure method of SDD vpath devices.

Table 8. Major files included in the SDD installation package (continued)

chgvpath	Method to change vpath attributes.
cfallvpath	Fast-path configuration method to configure the SDD pseudo-parent dpo and all vpath devices.
vpathdd	The SDD device driver.
hd2vp	The SDD script that converts an hdisk device volume group to a SDD vpath device volume group.
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.
querysn	The SDD driver tool to query unique serial numbers of devices.
mkvg4vp	The command that creates an SDD volume group.
extendvg4vp	The command that extends SDD 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 SDD devices.
restvg4vp	The command that restores all files belonging to a specified volume group with SDD 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 devices and LVM configuration when a migration failure occurs.
sddfcmap	Collects information on ESS SCSI or fibre-channel devices through SCSI commands.

Determining the installation type

Before you install 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 SDD installed on the host system, see "Installing SDD" on page 21 for instructions on installing and configuring SDD. If there is a previous version of 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

See "Migrating or upgrading SDD packages automatically without system restart" on page 32 for instructions on upgrading 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 35 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 SDD

SDD is released as an installation image. To install SDD, use the installation package that is appropriate for your environment.

Table 7 on page 17 lists and describes SDD support for 32-bit and 64-bit applications on AIX 4.3.3 or later.

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

If you are installing an *older* version of SDD, you must first remove any previously installed, *newer* version of SDD from your host system.

Installation of SDD software in a SAN File System environment might require special consideration. Consult the SAN File System documentation shown in "The SAN File System library" on page xxiii to check for any special steps that might need to be considered before proceeding with the SDD installation.

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).

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 proceeding with the manual upgrade instructions. See "Verifying if the SDD server has started" on page 48 and "Stopping the SDD server" on page 48 for more instructions. See "SDD server daemon" on page 48 for more details about the SDD server daemon.

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

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

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 your 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
```

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
```

Maximum number of LUNs

For different AIX OS levels and different types of devices, SDD has set different limits on the maximum number of LUNs that can be configured. These limits exist because AIX has limits on the number of devices that can be on a system. The combined number of hdisk and vpath devices should not exceed the number of devices that AIX supports. In a multipath environment, because each path to a disk creates an hdisk, the total number of disks being configured can be reduced by reducing the number of paths to each disk.

For versions 4.3 and 5.1, AIX has a published limit of 10 000 devices per system. Consequently, the maximum number of vpaths that can be configured is 600. This number is shared by all SDD supported devices. This means that if an AIX host has both ESS and SAN Volume Controller storage attached, or if an AIX host has both ESS and SAN Volume Controller for Cisco MDS 9000 storage attached, the two different types of devices share the maximum 600 as the total number of devices that SDD configures.

For version 5.2, AIX has removed the 10 000-device limit. Because of the increase in host system resources, SDD has increased the vpath limit accordingly. For AIX 5.2.0 (and later), you can have a configuration with a maximum of 1200 ESS LUNs and a maximum of 512 SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 LUNs. That is, different types of devices do not share the maximum number of LUNs limit on AIX 5.2.

Table 9 on page 25 provides a summary of the maximum number of LUNs allowed and the maximum number of paths allowed for a certain device when running on a particular host operating system. Because the number of paths might influence performance, you should use the minimum number of paths necessary to achieve redundancy in the SAN environment. The recommended number of paths is 2-4.

Note: The coexistence of SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 is not allowed.

Table 9. Maximum LUNs allowed for different AIX OS levels and different types of devices

OS level	ESS LUNs only	SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 LUNs only	ESS plus SAN Volume Controller LUNs	ESS plus SAN Volume Controller for Cisco MDS 9000 LUNs
AIX 4.3	Total of 600 LUNs maximum are allowed (maximum 32 paths)	Total of 600 LUNs maximum are allowed (maximum 32 paths)	Total of 600 LUNs maximum are allowed (maximum 32 paths)	Total of 600 LUNs maximum are allowed (maximum 32 paths)
AIX 5.1	Total of 600 LUNs maximum are allowed (maximum 32 paths)	Total of 600 LUNs maximum are allowed (maximum 32 paths)	Total of 600 LUNs maximum are allowed (maximum 32 paths)	Total of 600 LUNs maximum are allowed (maximum 32 paths)
AIX 5.2	1200 maximum LUNs are allowed (See Table 10 for maximum number of paths.)	512 maximum LUNs are allowed (maximum 32 paths)	Total of 1712 maximum LUNs (1200 maximum ESS LUNs + 512 SAN Volume Controller LUNs)	Total of 1712 maximum LUNs (1200 maximum ESS LUNs + 512 SAN Volume Controller for Cisco MDS 9000 LUNs)

For SDD 1.4.0.0 (or later), the maximum SDD device configuration for ESS LUNs on AIX 5.2 and the maximum paths per SDD device is given in Table 10.

Table 10. Maximum SDD device configuration for ESS LUNs on AIX 5.2

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.

Starting from SDD 1.4.0.0, three new SDD ODM attributes are available to control the maximum LUNs configuration:

- 2105_max_luns
- 2145_max_luns
- 2062_max_luns

The SDD ODM attribute, 2145_max_luns, defines the maximum number of SAN Volume Controller LUNs while 2062_max_luns defines the maximum number of SAN Volume Controller for Cisco MDS 9000 LUNs that can be configured on a host. Both attributes have a maximum value of 512. You cannot change these attributes.

You can have a maximum of 32 paths per vpath for SAN Volume Controller LUNs or SAN Volume Controller for Cisco MDS 9000 LUNs. Since the number of paths might influence performance, you should use the minimum number of paths necessary to achieve sufficient redundancy in the SAN environment.

The SDD ODM attribute, 2105_max_luns, defines the maximum number of ESS LUNs that can be configured on a host. This attribute is user changeable. The range of valid values for 2105_max_luns is 600 - 1200 in increments of 100. The default value for 2105_max_luns attribute is 600.

To display the values of the 2105_max_luns, 2145_max_luns, and 2062_max_luns attribute, use the command **lsattr -El dpo**:

```
> lsattr -El dpo
 2062_max_luns   512  Maximum LUNS allowed for 2062                 False
 2105_max_luns  1200  Maximum LUNS allowed for 2105                 True
 2145_max_luns   512  Maximum LUNS allowed for 2145                 False
 persistent_resv yes  Subsystem Supports Persistent Reserve Command False
 qdepth_enable   yes  Queue Depth Control                          True
```

The system administrator must ensure that the number of paths (hdisks) configured for ESS LUN does not exceed the maximum number of paths shown in Table 10 on page 25. If the number of paths (hdisks) for each LUN does exceed the maximum number of paths, the SDD configuration process terminates without configuring SDD devices.

This design is based on the current design of LVM **varyonvg** command. If there are hdisks that are not configured by SDD but sharing the same LUN as a vpath device, LVM might pick the hdisks physical volume instead of the SDD vpaths device during a SDD volume group varyon. This will cause the loss of single-point-failure protection provided by SDD.

The **addpaths** command follows the same path limitations for vpaths as identified in Table 10 on page 25. If **addpaths** finds that more paths (hdisks) are configured on a system than SDD allows, the command terminates.

Determining whether system has enough resource for increasing the maximum number of ESS LUNs

Before increasing the default value of SDD ODM attribute for ESS LUNs, 2105_max_luns, system administrators should first determine whether the system has sufficient resources for this operation.

ODM attributes: The AIX fibre-channel adapter has an ODM attribute named lg_term_dma that controls the DMA memory resource an adapter driver can use. When a host has more than 600 LUNs configured, the device open process might fail due to the lack of DMA memory resource. Before increasing the maximum number of LUNs, you should increase the lg_term_dma attribute. The default value of lg_term_dma is 0x200000 and the maximum value is 0x1000000. If you configure more than 600 LUNs, you should increase this attribute value to 0x400000 . If you still experience failure after changing this value to 0x400000, you should increase the value of this attribute again. Changing this attribute requires reconfiguration of the fibre-channel adapter and all its child devices. Because this is a disruptive procedure, You should change the lg_term_dma attribute before assigning or configuring ESS LUNs on a host system.

You should also change another fiber-channel adapter attribute, num_cmd_elems, which controls the maximum number of commands to be queued to the adapter. The default value is 200, whereas the maximum value is 2048. When a large number of ESS LUNs are configured, you can increase this attribute to improve performance.

Because reconfiguring a large number of devices is very time-consuming, you should perform the following steps to change the ODM attributes before configuring hdisks:

1. Execute **lsattr -El fcsN** to check the current value of lg_term_dma and num_cmd_elems.
2. Put all existing fibre-channel adapters and their children devices to the Defined state by issuing **rmdev -l fcsN -R**. It takes a long time to unconfigure a large number of devices. An alternative method, which can speed up this step, is to disconnect all fibre-channel cables and reboot the system.
3. Execute **chdev -l fcsN -a lg_term_dma=0x400000** to increase the DMA value.
4. Execute **chdev -l fcsN -a num_cmd_elems=1024** to increase the maximum commands value.
5. If you disconnected the fibre-channel cables in step 2, reconnect the cables.
6. Assign new LUNs to the AIX host.
7. Configure the fibre-channel adapters, the children devices and hdisks using **cfgmgr -l fcsN**.
8. With a large number of LUNs, many special device files will be created in the /dev directory. Executing the **ls** command with a wildcard (*) in this directory might fail. If executing the **ls** 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 (24K) and the maximum value for this attribute is 128 (512K). Increase the value of this attribute to 30. If you still experience failures after changing the value to 30, you should 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: After changing the ODM attributes to accommodate the increase of maximum number of LUNs, use the following steps to determine whether there is sufficient space in the root file system after hdisks are configured:

1. Execute **cfgmgr -l [scsiN/fcsN]** for each relevant SCSI or FCP adapter
2. Execute **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, execute **chfs** to increase the size of the root file system.

Increasing the maximum number of ESS LUNs

After installing SDD and preparing your system resource for more than 600 LUNs configuration, use the following procedures to configure more than 600 vpaths.

1. If vpaths are already configured on the system:
   a. Determine the current value of 2105_max_luns by using the command **lsattr -El dpo**. The following output is an example of the results of issuing the **lsattr -El dpo** command:

      ```
      > lsattr -El dpo
       2062_max_luns   512  Maximum LUNS allowed for 2062                False
       2105_max_luns  1200  Maximum LUNS allowed for 2105                True
       2145_max_luns   512  Maximum LUNS allowed for 2145                False
       persistent_resv yes  Subsystem Supports Persistent Reserve Command False
       qdepth_enable   yes  Queue Depth Control                          True
      ```

b. Determine how many paths each vpath currently has by issuing the command **datapath query device**.
c. Change all the existing vpaths state to DEFINED by executing **rmdev -l dpo -R**.
d. Ensure that the number of hdisks from the old configuration does not exceed the maximum paths allowed in the new configuration after you have increased the value of 2105_max_luns attribute. Otherwise, you need to remove the extra number of paths for each vpaths. See Table 10 on page 25 for maximum paths allowed.
e. Execute **/usr/lib/methods/defdpo**.
f. Execute **chdev -l dpo -a 2105_max_luns=***XXX*, where *XXX* is the maximum number of ESS LUNs that SDD can configure. Choose the value of 2105_max_luns from Table 10 on page 25.
g. Execute **cfallvpath** to configure SDD vpath devices.

2. If vpaths are not configured on the system, follow these steps:
   a. Determine that no more than the allowed number of hdisks are configured for each LUN. Otherwise, you need to remove the extra hdisks. For example, if the maximum LUN is intended between 901-1200, then no more than 4 hdisks per LUN should be configured.
   b. Execute **/usr/lib/methods/defdpo**.
   c. Execute **chdev -l dpo -a 2105_max_luns=***XXX*, where *XXX* is the maximum number of ESS LUNs that SDD can configure. Choose the value of 2105_max_luns from Table 10 on page 25.
   d. Execute **cfallvpath** to configure SDD vpath devices.

When configuring a large number of LUNs, you should enable fast failover to reduce the error recovery time. See “Enabling fast failover to reduce error recovery time” on page 52 for details

Preparing to configure SDD

Before you configure SDD, ensure that:

- The supported storage device is operational.
- The devices.sdd.*nn*.rte software is installed on the AIX host system, where *nn* 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 devices will lose multipath capability.

Perform the following steps:

1. Enter the **lsdev -Cc disk** command to check the hdisk configuration.

   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 -Cc disk | grep 2105** to check the ESS device configuration
   - **lsdev -Cc disk | grep SAN Volume Controller** to check the SAN volume controller device configuration
   - **lsdev -Cc disk | grep SVCCISCO** to check the SAN Volume Controller for Cisco MDS 9000 device configuration

2. If you have already created some ESS, SAN Volume Controller, or SAN Volume Controller for Cisco MDS 9000 volume groups, vary off (deactivate) all active volume groups with ESS, SAN Volume Controller, or SAN Volume Controller for Cisco MDS 9000 subsystem disks by using the **varyoffvg** (LVM) command. If some supported storage devices (hdisks) are used as physical volumes of an active volume group and if there are file systems of that volume group that are mounted, you must unmount all file systems and vary off all active volume groups with supported storage device SDD disks in order to configure SDD vpath devices correctly.

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

Starting with SDD 1.5.0.0, a new SDD attribute, qdepth_enable, allows you to control I/O flow to SDD vpath devices. 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 on your system, run the following command:

```
> lsattr -El dpo
  2062_max_luns   512  Maximum LUNS allowed for 2062                 False
  2105_max_luns  1200  Maximum LUNS allowed for 2105                 True
  2145_max_luns   512  Maximum LUNS allowed for 2145                 False
  persistent_resv yes  Subsystem Supports Persistent Reserve Command False
  qdepth_enable   yes  Queue Depth Control                          True
```

To disable queue depth logic, you should change the qdepth_enable attribute setting in ODM by executing **chdev**.

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 the **datapath set qdepth disable** command when the queue depth logic is currently enabled on the system, the following output is displayed:

```
+-----------------------------------------------------------------+
|Success: set qdepth_enable to no                                 |
+-----------------------------------------------------------------+
```

The SDD ODM attribute, qdepth_enable, will be updated. The following output is displayed when **lsattr -El dpo** is entered:

```
+-------------------------------------------------------------------------+
|persistent_resv yes Subsystem Supports Persistent Reserve Command False|
|qdepth_enable no Queue Depth Control True                                |
+-------------------------------------------------------------------------+
```

To disable queue depth logic for SDD versions earlier than 1.5.1.0, you should change the qdepth_enable attribute setting in ODM by executing **chdev**.

The following procedures are examples of changing the queue depth attribute under different SDD configuration conditions:

- If vpaths are already configured on the system, use the following procedure to change the value of the qdepth_enable attribute (in this case, from *yes* to *no*):
  1. Execute **rmdev -l dpo -R**.
  2. Execute **chdev -l dpo -a "qdepth_enable=no"**.
  3. Execute **lsattr -El dpo** to verify that the qdepth_enable attribute is changed. For example:

```
> lsattr -El dpo
2062_max_luns   512  Maximum LUNS allowed for 2062                False
2105_max_luns  1200  Maximum LUNS allowed for 2105                True
2145_max_luns   512  Maximum LUNS allowed for 2145                False
persistent_resv yes  Subsystem Supports Persistent Reserve Command False
qdepth_enable   no  Queue Depth Control                          True
```

  4. Execute **'cfallvpath'** to configure the vpaths.
- If vpaths are not configured on the system, to change the qdepth_enable value from *yes* to *no*:
  1. Execute **/usr/lib/methods/defdpo**.
  2. Execute **chdev -l dpo -a "qdepth_enable=no"**.
  3. Execute **lsattr -El dpo** to verify that the qdepth_enable attribute is changed.

```
> lsattr -El dpo
2062_max_luns   512  Maximum LUNS allowed for 2062                False
2105_max_luns  1200  Maximum LUNS allowed for 2105                True
2145_max_luns   512  Maximum LUNS allowed for 2145                False
persistent_resv yes  Subsystem Supports Persistent Reserve Command False
qdepth_enable   no  Queue Depth Control                          True
```

  4. Execute **'cfallvpath'** to configure the vpaths.

Note: These procedures are disruptive. If you are planning to run an application that will generate a large amount of I/O, you should perform the above procedures to disable the queue depth logic before you start the application.

Configuring SDD

Perform 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 vpath device configuration" on page 69.
5. Enter the **varyonvg** command to vary on all deactivated supported storage device volume groups.
6. If you want to convert the supported storage device hdisk volume group to SDD vpath devices, you must run the hd2vp utility. (See "hd2vp and vp2hd" on page 85 for information about this utility.)
7. 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.
2. Run the vp2hd conversion script to convert the volume group from SDD devices (vpath*N*) to supported storage devices (hdisks).

Note: With SDD 1.3.3.11 (or earlier), if you are running HACMP with the ibmSdd_433.rte or ibmSdd_510nchacmp.rte installation package installed on your host system, there are special requirements regarding unconfiguring and removing SDD vpath devices.

With SDD 1.4.0.0 (or later), if you are running HACMP with the devices.sdd.43.rte or devices.sdd.51.rte installation package installed on your host system, there are special requirements regarding unconfiguring and removing SDD vpath devices.

See “Special requirements” on page 58.

You can use SMIT to unconfigure the SDD devices in two ways. Either you can 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 vpath***N* 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 30.

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

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 **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 or not 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**.

Migrating or upgrading 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 numerous volume groups and file systems are created with SDD devices. This feature will support:

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.

During the migration or upgrade of 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 before running the migration:

1. If HACMP is running, gracefully stop the cluster services.
2. If sddServer.rte (standalone 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 SDD migration or upgrade (for example, configuring more devices).
5. If there is active paging space created with SDD devices, deactivate the paging space.
6. SDD does not support mixed volume groups with SDD devices and non-SDD devices. A volume group should contain SDD devices only or non-SDD devices only. If the non-SDD device is hdisk, to fix the mixed volume group to contain hdisks only, execute the following command:

   vp2hd <*volume group name***>**

   To fix the mixed volume group to contain SDD devices only, simply start the SDD migration or upgrade and the mixed volume group will be fixed automatically by the migration scripts. The following messages will be displayed during the migration or upgrade:

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

   .
   ```

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

Procedures for migration or upgrade

To start SDD migration or upgrade:

1. Install new SDD package by entering **smitty install**. The migration or upgrade script to save the LVM configuration will be executed as part of the installation procedure initiated by **smitty install**.

   You will see the following messages indicating that the pre-uninstallation operations of 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 will be uninstalled before new SDD will be installed.
3. The migration or upgrade script will automatically configure SDD devices and restore the LVM configuration.

   You will see the following messages indicating that the post-installation of SDD is successful:

   ```
   Original lvm configuration is restored successfully!
   ```

Error recovery for migration or upgrade

If any error occurred during the pre-installation or post-installation procedures, such as disconnection of cables, you can recover the migration or upgrade. Two common ways that migration or upgrade can fail are:

Case 1: **Smitty install** failed.

Smitty install will fail if there is an error during the pre-uninstallation activities for the older SDD package. An error message indicating the error will be printed. You should identify and fix the problem. Then 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 post-installation (either the configuration of SDD devices failed or LVM restoration failed), the new SDD package will still be successfully installed. Thus, the **Smitty install** will exit with an OK prompt. However, an error message indicating the error will be printed. You should identify and fix the problem. Then, run the shell script **lvmrecover** to configure SDD devices and automatically recover the original LVM configuration.

Migrating or upgrading SDD manually

The following section describes the procedure to migrate or upgrade SDD manually. See “Migrating or upgrading SDD packages automatically without system restart” on page 32 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 11 on page 34.
- Upgrading the AIX operating system and thus upgrading SDD package. For example, upgrading AIX 4.3 to AIX 5.1.

You must uninstall the existing SDD and then manually install the new version of SDD in these cases.

Table 11. List of previously installed installation packages that are supported with the installation upgrade

Installation package name
ibmSdd_432.rte
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

Perform the following steps to upgrade SDD:

1. Remove any .toc files generated during previous SDD installations. Enter the following command to delete any .toc file found in the /usr/sys/inst.images directory:

   rm .toc

   Note: Ensure that this file is removed because it contains information about the previous version of SDD.
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. Run the **vp2hd** script to convert the volume group from SDD devices to ESS hdisk devices.
7. Remove all SDD devices. Enter the following command:

   ```
   rmdev  -dl dpo -R
   ```
8. Use the **smitty** command to uninstall 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 39 for the step-by-step procedure for uninstalling 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 SDD from the compact disc. Enter **smitty install** and press **Enter**. The installation process begins. Go to “Installing SDD” on page 21 to complete the installation process.
11. Use the **smitty device** command to configure all the SDD devices to the *Available* state. See “Configuring SDD” on page 30 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. Run the **hd2vp** script for each SDD volume group to convert the physical volumes from ESS hdisk devices back to SDD vpath devices. Enter the following command:

    ```
    hd2vp  vg_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 4 on page 34. 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 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.

Migrating or upgrading the SDD package during an AIX OS upgrade

SDD provides different packages to match the AIX OS level. If an AIX system is going to be upgraded to a different OS level, then you need to install the corresponding SDD package for that OS level. Automatic migration of an SDD package from a lower OS level to a later OS level after an OS upgrade is not supported. For example, migration from devices.sdd.43.rte to devices.sdd.51.rte after an OS upgrade from AIX 4.3 to AIX 5.1, or 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. See "Migrating or upgrading SDD manually" on page 33 for information about manually migrating or upgrading SDD.

Important:

1. The maximum number of supported LUNs for the SAN Volume Controller on AIX 4.3.3 and AIX 5.1 is 600, whereas the maximum number of supported LUNs for the SAN Volume Controller on AIX 5.2 is 512.
2. The maximum number of supported LUNs for the SAN Volume Controller for Cisco MDS 9000 on AIX 5.1 is 600, whereas the maximum number of supported LUNs for the SAN Volume Controller for Cisco MDS 9000 on AIX 5.2 is 512.

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. The SMIT facility has two interfaces, nongraphical (enter **smitty** to invoke the nongraphical user interface) and graphical (enter **smit** to invoke the GUI).

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

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

Perform 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 perform 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. Run the **vp2hd** script to convert the volume group from SDD devices to ESS or SAN Volume Controller hdisk devices.
7. Remove all SDD devices. Enter the following command:

   ```
   rmdev -dl dpo -R
   ```

Perform 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).

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

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 devices to the Available state. See “Configuring fibre-channel-attached devices” on page 15 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. Run the **hd2vp** script for each SDD volume group to convert the physical volumes from ESS or SAN Volume Controller 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 will not function properly. Enter the **dpovgfix vg_name** command to fix this problem.

7. Start SDD Server. Enter the following command:

   ```
   startsrc -s sddsrv
   ```

Verifying the SDD configuration

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

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

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 **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 pseudo 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 30 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 pseudo device (for example, vpath13)
- The *Defined* or *Available* state of a pseudo device
- Whether or not the pseudo 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 ESS LUN (for example, 02FFA067) or the unit serial number of the SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 LUN (for example, 60056768018A0210B00000000000006B)
- The names of the AIX disk devices making up the pseudo device and their configuration and physical volume state

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 48 and "Stopping the SDD server" on page 48 for more instructions. See Chapter 11, "Using the SDD server and the SDDPCM server," on page 279 for more details about the SDD server daemon.

Before you remove the SDD package from your AIX host system, all the SDD devices must be unconfigured and removed from your host system. See "Unconfiguring SDD" on page 31.

The fast-path **rmdev -dl dpo -R** command removes all the SDD devices from your system. After all SDD devices are removed, perform the following steps to remove SDD.

1. Enter **smitty deinstall** from your desktop window to go directly to the Remove Installed Software panel.
2. 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

   Then press **Enter**.

   Note: See "Verifying the currently installed version of SDD for SDD 1.3.3.11 (or earlier)" on page 22 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.

3. 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 process stops at this point and previews what you are removing. The results of your pre-check are displayed without removing the software. If the state for any SDD 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 SDD software package is removed from your system.

Preferred node path-selection algorithm for SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000

The SAN Volume Controller and the SAN Volume Controller for Cisco MDS 9000 is a two-controller disk subsystem. SDD distinguishes the paths to a SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 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 on the preferred controller can be used 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 I/O retries fail on all paths on the preferred controller, then 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 paths on the alternate controller.
3. After failover to the alternate controller has occurred, if a path on the preferred controller becomes available, I/O will automatically fail back to the paths on the preferred controller.

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*.

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.

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 vpath***N* command, where *N* represents the vpath number. Press **Enter**. The output should look similar to this:

```
pvid         0004379001b90b3f0000000000000000 Data Path Optimizer Parent False
policy       df                               Scheduling Policy          True
active_hdisk hdisk1/30C12028                  Active hdisk               False
active_hdisk hdisk5/30C12028
serial_number 01523922 N/A 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** *rr/fo/lb/df* command to change the policy dynamically associated with vpaths in either *Close* or *Open* state.

See “datapath set device policy” on page 303 for more information about the **datapath set device policy** command.

Dynamically adding paths to SDD devices of a volume group

With SDD 1.3.1.3 (or later), you can dynamically add more paths to SDD devices after you have initially configured SDD. This section shows you how to add paths to SDD devices from AIX 4.3.2 (or later) host systems with the **addpaths** command.

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

If you enter the **addpaths** command to a vpath device that is in the Open state, 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 a 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 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 will be automatically in the Open state after you enter the **addpaths** command to a vpath device that has only one configured path.

There are special considerations if you are using **addpaths** with AIX 5.2.0. SDD limits the number of paths per vpath that you can have with AIX 5.2.0, depending on the number of vpaths you have.

You can determine how many more hdisks can be added to the existing SDD devices by using the commands **lsattr -El dpo** and **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 maximum paths allowed according to the value of 2105_max_luns. See Table 10 on page 25 for the maximum paths allowed.

For example, on an AIX 5.2.0 host that has 2105_max_luns set to a value of 900, if there are four hdisks configured to each vpath, only four more hdisks per vpath can be added. If there are eight hdisks configured to each vpath, no more hdisks can be added to the vpaths.

Note to system administrators: If you configure more hdisks than are allowed, running **addpaths** will not add any paths to vpaths.

Aside from the previous special consideration for AIX 5.2.0 systems, the procedure for adding paths to SDD devices for all supported OS levels are the same.

Before you issue the **addpaths** command, make sure that ESS logical volume sharing is enabled for all applicable devices. You can enable ESS logical volume sharing through the ESS Specialist. Refer to the *IBM TotalStorage Enterprise Storage Server: Web Interface User's Guide* for more information.

If you would like to add paths to SDD devices which are not a volume group, you can execute the **cfgmgr** command *n* times, where *n* represents the number of paths for SDD, and then issue the **addpaths** command from the AIX command line to add more paths to the SDD devices. If you would like to add paths to SDD devices of a volume group, complete the following steps:

1. Enter the **lspv** command to list the physical volumes.
2. Identify the volume group that contains the SDD devices to which you want to add more paths.
3. Verify that all the physical volumes belonging to the SDD volume group are SDD devices (vpath*N*s). If they are not, you must fix the problem before proceeding to the next step. Otherwise, the entire volume group loses the path-failover protection.

   You can issue the **dpovgfix vg-name** command to ensure that all physical volumes within the SDD volume group are SDD devices.
4. Run the AIX configuration manager to recognize all new hdisk devices. Ensure that all logical drives on the supported storage device are identified as hdisks before continuing. Select one way to run the AIX configuration manager:

- Enter the **cfgmgr** command *n* times, where *n* represents the number of paths for SDD
- Enter the **cfgmgr -l [scsiN/fcsN]** command for each relevant SCSI or FCP adapter.

5. Enter the **addpaths** command from the AIX command line to add more paths to the SDD devices.
6. Enter the **lsvpcfg** command from the AIX command line to verify the configuration of the SDD devices in the volume group.

   SDD devices should show two or more hdisks associated with each SDD device when the failover protection is required.

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 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 286 in Chapter 12, “Using the datapath commands,” on page 283 for more information.

Dynamically replacing adapters or removing paths

SDD v1.5.1.0 (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 a 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 replacing another adapter of the same type”
- “Dynamically replacing an adapter with an adapter of a different type” on page 44
- “Dynamically removing a path of a SDD vpath device” on page 46

Dynamically replacing another adapter of the same type

To dynamically replace an adapter of the same type, use the **datapath set adapter n suspend** command. This command sets the state of the adapter and its associated paths to SUSPEND and their mode to OFFLINE. All suspended paths are closed and I/O will not be routed to these paths. See “datapath set adapter” on page 301 for more information about the **datapath set adapter n suspend** command.

After you physically replace the existing adapter with a new one of the same type, the new adapter and its child devices must be configured to Available state. Then, use the **datapath set adapter n resume** command to reconfigure the adapter and its child devices to SDD. All logical device names associated with the previous adapter in SDD are reused for the new one. Upon a successful execution, the previously suspended adapter and its associated paths are resumed with the adapter state set to NORMAL and its mode set to ACTIVE. The state of paths is set to either OPEN or CLOSE depending on the state of vpath to which the paths belong. The path mode is set to NORMAL. See “datapath set adapter” on page 301 for more information about the **datapath set adapter n resume** command.

Complete the following steps to dynamically replace an adapter of the same type in the SDD environment:

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        425     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        |
+-------------------------------------------------------------------------------------+
```

2. Enter **datapath set adapter** *n* **suspend**, where *n* is the adapter number to be replaced as shown in the output of the **datapath query adapter** command. For example, to replace Adpt# 0, enter **datapath set adapter 0 suspend**.

```
+-------------------------------------------------------------------------------------+
|Success: set adapter 0 to suspend                                                    |
|                                                                                     |
|Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active        |
|    0           fscsi0 SUSPEND  OFFLINE      62070        425     10      10        |
+-------------------------------------------------------------------------------------+
```

3. Enter **rmdev -dl fcs0 -R** to remove fcs0, a parent of fscsi0, and all of its child devices from the system. Executing **lsdev -Ccdisk** should not show any devices associated with fscsi0.
4. Enter **drslot -R -c pci -s P1-I8** where P1-I8 is the slot location found from executing **lscfg -vl fcs0**. This command prepares a hot-plug slot for systems with AIX 5L or later.
5. Follow the instructions of the **drslot** command to physically remove the adapter and install a new one. Remember that the new adapter must be the exact same type of adapter as the old one.
6. Update the WWN of the new adapter at the device end and in the fabric. For ESS, 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.
7. Enter **cfgmgr** or **cfgmgr -vl pci(***n***)**, where *n* is the pci number, to configure the new adapter and its child devices. Use the **lsdev -Cc hdisk** and **lsdev -Cc adapter** commands to ensure that all devices are successfully configured to *Available* state.
8. Enter **datapath set adapter** *n* **resume**, where *n* is the adapter number, to resume the newly installed adapter and its child devices for SDD. Resumed paths will be opened automatically if vpath is open.

```
+-------------------------------------------------------------------------------------+
|Active Adapters :4                                                                   |
|                                                                                     |
|Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active        |
|    0           fscsi0  NORMAL   ACTIVE         56          0     10      10        |
|    1           fscsi1  NORMAL   ACTIVE     186686         10     10      10        |
|    2           fscsi2  NORMAL   ACTIVE     107677         54     10      10        |
|    3           fscsi3  NORMAL   ACTIVE     104999         46     10      10        |
+-------------------------------------------------------------------------------------+
```

Dynamically replacing an adapter with an adapter of a different type

To dynamically replace an adapter of a different type (for example, an adapter with a different speed) or to permanently remove an adapter and its child devices from SDD, use the **datapath remove adapter** *n* command, where *n* is the adapter number. In this case, all logical device names associated with the adapter are deleted from SDD. Upon a successful execution, this adapter and associated paths will not occur in any **datapath query** or **lsvpcfg** command output.

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 298 for more information about the **datapath remove adapter** *n* command.

Complete the following steps to dynamically replace an adapter of a different type in the SDD environment:

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     |
+-------------------------------------------------------------------------------------+
```

2. 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.

3. Enter **rmdev -dl fcs0 -R** to remove fcs0, a parent of fscsi0, and all of its child devices from the system. Executing **lsdev -Cc disk** should not show any devices associated with fscsi0.
4. Enter **drslot -R -c pci -s P1-I8** where P1-I8 is the slot location found by executing **lscfg -vl fcs0**. This command prepares a hot plug slot for systems with AIX 5L or later.
5. Follow the instruction given by **drslot** to physically remove the adapter and install a new one. Remember that the new adapter can be a different type from the old one.
6. Update the WWN of the new adapter at the device end and in the fabric. For ESS, 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.
7. 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 hdisk** and **lsdev -Cc adapter** commands to ensure that all devices are successfully configured to *Available* state.
8. 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 a SDD vpath device

To dynamically remove a particular path from a SDD vpath device, use the **datapath remove device path** command. In this case, the particular logical path is permanently deleted from SDD. Upon a successful execution, this path will not be shown in any **datapath query** or **lsvpcfg** command output. See “datapath remove device path” on page 299 for more information about the **datapath remove device path** command.

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

1. 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 DEV# 0 Path # 1.

```
+-----------------------------------------------------------------------------------------+
|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. Enter **addpaths** to reclaim the removed path.

```
+-----------------------------------------------------------------------------------------+
|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.

Fibre-channel Dynamic Device Tracking for AIX 5.20 ML1

This section applies only to the devices.sdd.52.rte package on AIX 5.20 ML1.

With AIX 5.20 ML1 and later releases, the AIX fibre-channel driver will support fibre-channel Dynamic Device Tracking. This will enable the dynamic changing of fibre-channel cable connections on switch ports or on supported storage ports

without unconfiguring and reconfiguring hdisk and vpath devices. SDD 1.5.0.0 and later will support this new feature and allow 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.
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.
4. Swap ports of two cables on ESS; the disconnected cable must be reconnected within 15 seconds.

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

1. Execute the **rmdev -l fscsiX -R** for all adapters on your system to change all the children devices of fscsiX on your system to the defined state.
2. Execute 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 **lslpp** 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 ESS where a cable is moved from one adapter to another free, previously unseen adapter on the ESS. 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.

Using ESS concurrent download of licensed internal code

Concurrent download of licensed internal code (LIC) is the capability to download and install licensed internal code on an ESS while applications continue to run. This capability is supported for multiple-path (SCSI or FC) access to an ESS.

Attention: You ***should not*** shut down the host during a concurrent download (CCL) of licensed internal code.

For information about performing concurrent download of LIC for the ESS, refer to the Microcode Installation Instructions for your specific type and model.

SDD server daemon

The SDD server (also referred to as 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 11, “Using the SDD server and the SDDPCM server,” on page 279 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 49.

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.

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

See “Changing the sddsrv or pcmsrv TCP/IP port number” on page 282.

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 rebooted.

You can also choose to replace the current version of sddsrv with a standalone version by doing the following:

1. Enter **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 48 and the status should be *inoperative*.
3. Comment out following line:

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

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

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

   to the system /etc/inittab table.
5. Enter **startsrc -a s0 -s sddsrv** to start a standalone 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 304 for more information.

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 304 for more information.

If you are running HACMP and are experiencing problems associated with sddsrv (see “Understanding SDD support for High Availability Cluster Multi-Processing” on page 53), see Table 12 for information about HACMP fixes that will solve the problems.

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 sddsrv until the PTFs for these APARs can be installed. After the PTFs for these APARs are installed, you should re-enable sddsrv.

Understanding SDD 1.3.2.9 (or later) support for single-path configuration for ESS

SDD 1.3.2.9 (or later) does not support concurrent download and installation of the licensed internal code (LIC) on single-path mode.

However, SDD supports single-path SCSI or fibre-channel connection from your AIX host system to an ESS. It is possible to create a volume group or a vpath device with only a single path.

Note: Because of single-path connection, SDD can not provide single-point-failure protection and load balancing. IBM does not recommend this.

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 devices of a volume group” on page 41.

Understanding SDD error recovery policies

There are differences in the way that SDD 1.3.2.9 (or later) and SDD 1.3.1.3 (or earlier) handle error recovery for AIX host systems:

SDD 1.3.1.3 (or earlier) error recovery policy
With SDD 1.3.1.3 (or earlier), the error recovery policy is designed to withstand transient errors from user applications. The error recovery policy prevents a path from becoming disabled in the event of transient errors occurring. This policy can cause error recovery to take a long time. The recovery policy halts the I/O activities on functional paths and vpath devices for some time before failing paths are set to the DEAD state.

SDD 1.3.2.9 (or later) error recovery policy
With SDD 1.3.2.9 (or later), error recovery policy is designed to report failed I/O requests to applications more quickly. This process prevents unnecessary retries, which can cause the I/O activities on good paths and vpath devices to halt for an unacceptable period of time.

Both the SDD 1.3.2.9 (or later) and SDD 1.3.1.3 (or earlier) error recovery policies support the following modes of operation:

single-path mode (for ESS only)
An AIX host system has only one path that is configured to an ESS LUN.

SDD, in single-path mode, has the following characteristics:

- When an I/O error occurs, SDD retries the I/O operation up to two times.
- With the SDD 1.3.2.9 (or later) error recovery policy, SDD returns the failed I/O to the application and sets the state of this failing path to DEAD. SDD driver relies on the SDD server daemon to detect the recovery of the single path. The SDD server daemon recovers this failing path and changes its state to OPEN. (SDD can change a single and failing path into DEAD state.)

- With the SDD 1.3.1.3 (or earlier) error recovery policy, SDD returns the failed I/O to the application and leaves this path in OPEN state (SDD never puts this single path into DEAD state).
- With SDD 1.3.2.9 (or later), the SDD server daemon detects the single CLOSE path that is failing and changes the state of this failing path to CLOSE_DEAD. When the SDD server daemon detects that a CLOSE_DEAD path has recovered, it will change the state of this path to CLOSE. With a single path configured, the SDD vpath device cannot be opened if the only path is in a CLOSE_DEAD state.

multipath mode

The host system has multiple paths that are configured to a supported storage device LUN.

Both the SDD 1.3.2.9 (or later) and SDD 1.3.1.3 (or earlier) error recovery policies in multipath mode have the following common characteristics:

- If an I/O error occurs on the last operational path to a device, SDD attempts to reuse (or fail back to) a previously-failed path.

The SDD 1.3.2.9 (or later) error recovery policy in multiple-path mode has the following latest characteristics:

- If an I/O error occurs on a path, SDD 1.3.2.9 (or later) does not attempt to use the path until three successful I/O operations occur on an operational path.
- If an I/O error occurs consecutively on a path and the I/O error count reaches three, SDD immediately changes the state of the failing path to DEAD.
- Both SDD driver and the SDD server daemon can put a last path into DEAD state, if this path is no longer functional. The SDD server can automatically change the state of this path to OPEN after it is recovered. Alternatively, you can manually change the state of the path to OPEN after it is recovered by using the **datapath set path online** command. Go to “datapath set device path” on page 304 for more information.
- If the SDD server daemon detects that the last CLOSE path is not functional, the daemon will change the state of this path to CLOSE_DEAD. The SDD server can automatically recover the path if it is determined to be functional.
- If an I/O fails on all OPEN paths to a storage device LUN, SDD returns the failed I/O to the application and changes the state of all OPEN paths (for failed I/Os) to DEAD, even if some paths might not reach I/O error count to three.
- If an OPEN path already failed some I/Os, it will not be selected as a retry path.

The SDD 1.3.1.3 (or earlier) error recovery policy in multipath mode has the following characteristics:

- If an I/O error occurs on a path, SDD 1.3.1.3 does not attempt to use the path until 2000 successful I/O operations occur on an operational path.
- The last path is reserved in OPEN state.
- If an I/O fails on all OPEN paths to an ESS LUN, SDD returns the failed I/O to the application and leaves all the paths in OPEN state.
- A failed I/O is retried on all OPEN paths to an ESS LUN even if the OPEN path already failed I/Os.

- SDD changes the failed path from the DEAD state back to the OPEN state after 50 000 successful I/O operations on an operational path.

Enabling fast failover to reduce error recovery time

In AIX 5.1 and AIX 5.2B, the fc_err_recov attribute enables fast failover during error recovery. Enabling this attribute can reduce the amount of time SDD needs to fail a broken path. The default value for fc_err_recov is delayed_fail.

Notes:

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).

To enable fast failover, do the following:

1. Change all the children devices of fscsiX on your system to the defined state by executing **'rmdev -l fscsiX -R'** for all adapters on your system.
2. Execute the **'chdev -l fscsiX -a fc_err_recov=fast_fail'** command for all adapters on your system.
3. Run cfgmgr to reconfigure all devices back to the available state.

Understanding SDD support for pSeries 690 with static LPARs configured

The pSeries 690 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 pSeries hardware resources with any other partition. So each partition functions the same way that it does on a standalone 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, then the path optimization can be performed on those adapters in the same way as in a standalone 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).

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. Execute **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 executing **lquerypr -Vrh /dev/vpathX** command. After successful execution, you will be able to vary on the volume group successfully.

Understanding SDD support for High Availability Cluster Multi-Processing

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 Multi-Processing (HACMP) running on RS/6000 and pSeries 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 (corresponding 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, IBM recommends that you 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, refer to the corresponding User's Guide for HACMP support information. See Table 13.

HACMP provides a reliable way for clustered IBM RS/6000 and pSeries servers that share disk resources to recover from server and disk failures. In an HACMP environment, each RS/6000 or pSeries 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 13. Recommended SDD installation packages and supported HACMP modes

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

Tip: If 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 if you are running an SDD version earlier than 1.4.0.0.

HACMP is not supported on all models of the ESS. For information about supported ESS models and required ESS microcode levels, go to the following Web site:

www.ibm.com/servers/storage/support/software/sdd.html

SDD supports RS/6000 and pSeries 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.

The following tables summarize the software requirements to support HACMP:

- Table 14
- Table 15 on page 55
- Table 16 on page 56
- Table 17 on page 57

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

Table 14. Software support for HACMP 4.4, 4.4.1, and 4.5 on AIX 4.3.3 (32-bit only), 5.1.0 (32-bit and 64-bit), and 5.2.0 (32-bit and 64-bit)

SDD version and release level	HACMP 4.4 + APARs	HACMP 4.4.1 + APARs	HACMP 4.5 + APARs
devices.sdd.43.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)	• IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • IY14682 • F model requires IY11110	• IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • IY14683 • F model requires IY11480	Not applicable
devices.sdd.51.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)	• IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480	• IY26561 • (See note.)	• IY36938 • IY36933 • IY35735 • IY36951

Table 14. Software support for HACMP 4.4, 4.4.1, and 4.5 on AIX 4.3.3 (32-bit only), 5.1.0 (32-bit and 64-bit), and 5.2.0 (32-bit and 64-bit) (continued)

SDD version and release level	HACMP 4.4 + APARs	HACMP 4.4.1 + APARs	HACMP 4.5 + APARs
devices.sdd.52.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)	Not applicable	Not applicable	• IY36938 • IY36933 • IY36782 • IY37744 • IY37746 • IY35810 • IY36951
Note: For up-to-date APAR information for HACMP, go to the following Web site: https://techsupport.services.ibm.com/server/aix.fdc			

Table 15. Software support for HACMP 4.4, 4.4.4, and 4.5 on AIX 5.1.0 (32-bit and 64-bit kernel)

SDD version and release level	HACMP 4.4 + APARs	HACMP 4.4.1 + APARs (64-bit)	HACMP 4.5 + APARs
ibmSdd_510nchacmp.rte installation package for SDD 1.3.1.3 (SCSI/FCP)	• IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480	• IY26561 • (See note.)	• IY36938 • IY36933 • IY35735 • IY36951
ibmSdd_510nchacmp.rte installation package for SDD 1.3.2.9 (to SDD 1.3.3.x) (SCSI/FCP)	• IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480	• IY26561 • IY25995 • (See note.)	• IY36938 • IY36933 • IY35735 • IY36951
Note: Only HACMP 4.4.1 and APARs support 64–bit mode. For up-to-date APAR information for HACMP, go to the following Web site: https://techsupport.services.ibm.com/server/aix.fdc			

Table 16. Software support for HACMP 4.3.1 and 4.4 in concurrent mode on AIX 4.3.3

SDD version and release level	HACMP 4.3.1 + APARs	HACMP 4.4 + APARs	HACMP 4.4.1+ APARs
ibmSdd_432.rte installation package for SDD 1.1.4.0 (SCSI only)	• IY07392 • IY03438 • IY11560 • IY08933 • IY11564 • IY12021 • IY12056 • F model requires IY11110	• IY11563 • IY11565 • IY12022 • IY12057 • F model requires IY11480	Not applicable.
ibmSdd_432.rte installation package for SDD 1.2.0.0 and 1.2.2.X (SCSI/FCP)	• IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • F model requires IY11110	• IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • F model requires IY11480	Not applicable.
ibmSdd_432.rte installation package for SDD 1.3.1.3 (to SDD 1.3.3.x) (SCSI/FCP)	• IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • F model requires IY11110	• IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • F model requires IY11480	(See note.)
Note: For up-to-date APAR information for HACMP, go to the following Web site: https://techsupport.services.ibm.com/server/aix.fdc			

Table 17. Software support for HACMP 4.3.1 and 4.4 in nonconcurrent mode on AIX 4.3.3

SDD version and release level	HACMP 4.3.1 + APARs	HACMP 4.4 + APARs	HACMP 4.4.1 + APARs
ibmSdd_433.rte installation package for SDD 1.2.2.X (SCSI/FCP)	• IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • IY14682 • F model requires IY11110	• IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • IY14683 • F model requires IY11480	Not applicable.
ibmSdd_433.rte installation package for SDD 1.3.1.3 (to SDD 1.3.3.x) (SCSI/FCP)	• IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • IY14682 • F model requires IY11110	• IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • IY14683 • F model requires IY11480	(See note.)
Note: For up-to-date APAR information for HACMP, go to the following Web site: https://techsupport.services.ibm.com/server/aix.fdc			

HACMP and sddsrv

The SDD server daemon 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 SDD. The SDD server daemon starts automatically after the SDD package is installed. It scans to find failing paths regularly at 2- to 5-minutes intervals. It probes the idle paths at 1-minute internals.

During the HACMP fallover process, the fallover time can take longer to complete when running with broken paths if the SDD server daemon is running in the background. This problem is caused by the current locking design in AIX fibre-channel driver. If you experience this problem, you should install the fixes listed in Table 18. When the fixes are installed, sddsrv can be re-enabled. If you have removed or commented out the line in the system /etc/inittab file that automatically starts sddsrv after a system reboot, you should replace or remove the comment notation from that line in the system /etc/inittab file.

Table 18. APARs on AIX with HACMP and the SDD server daemon running

HACMP version	APAR
HACMP v4.4.1	IY36443
HACMP v4.5	IY35382

SDD persistent reserve attributes

With SDD 1.4.0.0 or later, a single package (corresponding 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. They require the ESS G3-level microcode on the ESS to support the SCSI-3 Persistent Reserve command set. If the ESS G3-level microcode is not installed, the SDD installation packages will default to single-path configurations instead of multipath configurations. There is no single-point-failure protection for single-path configurations.

The SDD installation packages have a new attribute under the pseudo-parent (dpo) that reflects whether or not the supported storage device supports the Persistent Reserve Command set. The attribute name is *persistent_resv*. If SDD detects that G3-level microcode is installed, 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 with the G3 microcode, the output should look similar to this:

```
persistent_resv yes N/A False
```

In order to implement the persistent reserve command set, each host server needs a unique, 8-byte reservation key. You can get a unique reservation key in one of two ways. In HACMP environments, HACMP generates a unique key for each node in the ODM database. In case HACMP is not installed on an AIX server, SDD configure methods generate a unique reservation key by using the six digits of the machine ID from issuing the **uname -m** command and appending two digits of 00 to the machine ID. With an AIX server supporting LPAR configuration (such as a pSeries 690 system), if HACMP is not installed and SDD persistent reserve version is installed, the SDD configuration methods generate a unique reservation key by appending two digits of LPAR ID from the output of the **uname -L** command to the six digits of the machine ID from issuing the **uname -m** command.

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
```

Special requirements

There is a special requirement regarding unconfiguring and removing SDD vpath devices when running in a HACMP environment. You must unconfigure and remove the vpath devices *before* you unconfigure and remove the underlying supported storage device hdisks for the vpath devices. Otherwise, if the supported storage

device hdisks are unconfigured and removed first, the persistent reserve will not be released on the physical devices, even though the vpath devices have been successfully unconfigured and removed.

SDD does not automatically create the pvid attribute in the ODM database for each 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 does not automatically create the pvid attribute in the ODM database for each 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 85) to convert the volume group's physical volumes from supported storage device hdisks to vpath devices. This conversion step not only creates pvid attributes for all vpath devices that belong to that imported volume group, it also deletes the pvid attributes for the underlying hdisks for these vpath devices. Later on, you can import and vary on the volume group directly from the 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 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 75 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 AIX v5.1.D and HACMP v4.4.1.4, 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 64.

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

HACMP RAID concurrent-mode volume groups

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

Creating HACMP RAID concurrent-mode volume groups

Perform 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 the entry fields.
Press Enter AFTER making all required changes.

 [Entry Fields]
VOLUME GROUP name               [con_vg]
Physical partition SIZE in megabytes           4
PHYSICAL VOLUME names             [vpath124]
Activate volume group AUTOMATICALLY at system restart? no
Volume Group MAJOR NUMBER            [80]
Create VOLUME GROUPS concurrent-capable? no
Auto-varyon in concurrent mode? no
LTG size in kbytes 128
*******************************************************************************
```

Importing HACMP RAID concurrent-mode volume groups

When importing the volume group to other nodes in the cluster, the user must import the volume group to the hdisks and then manually move the volume group definition to the vpath. The hdisk values can be obtained in a number of ways:

Attention: Before you begin to import to other nodes, you need to vary off the volume group on the node after it is created.

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:

   ```
   NODE VG BEING IMPORTED TO
     zebra> lspv | grep 000900cf4939f79c
   hdisk126  000900cf4939f79c con_vg
   hdisk252  000900cf4939f79c con_vg
   zebra>
   zebra> lsvpcfg
   vpath124 (Avail) 21B21411=hdisk126 (Avail pv) hdisk252 (Avail pv)
   ```

   Notes:

   a. If the hdisks do not have pvid, issue the **chdev -l hdiskX -a pv=yes** command to recover it. Verify that the pvid exists by issuing the **lspv** command against the hdisk. Use the same method when attempting to obtain the pvid on vpath.

   b. 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.

The following procedure will be helpful in determining which hdisk you need to issue the **chdev -l hdiskX -a pv=yes** commands on.

Hdisks can be obtained by issuing the **lsvpcfg** command.

```
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>
```

Example 1:

```
zebra> lsvpcfg
vpath122 (Avail) 21921411=hdisk255 (Avail) hdisk259 (Avail)
vpath123 (Avail) 21A21411=hdisk256 (Avail) hdisk260 (Avail)
vpath124 (Avail) 21B21411=hdisk126 (Avail pv) hdisk252 (Avail pv)
zebra>
```

Note: There is a possibility that the vpath value might be different for each host. Verify that the value (21B21411) following the vpath name on each node are identical.

After you have verified the values following the vpath name, you need to perform the following steps.

1. Enter **smitty importvg** at the command prompt.

   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
********************************************************************************
```

2. After importing volume groups has been completed, issue the **lsvpcfg** command again to verify the state of the vpath.

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

3. Enter the **hd2vp** command against the volume group:

```
zebra> hd2vp con_vg
zebra> lsvpcfg
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.

Removing HACMP RAID concurrent-mode volume groups

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

Notes:

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.

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 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

Perform 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 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

Perform 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 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

Perform the following step to export an HACMP RAID concurrent-mode volume group.

1. Enter **exportvg** at the command prompt.

   Notes:

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

Enhanced concurrent-capable volume groups

With the AIX v5.1.D and HACMP v4.4.1.4 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.

Attention: There is a defect in the AIX LVM driver associated with extending physical volumes of an enhanced concurrent volume group on more

than two nodes in an HACMP environment. The **extendvg4vp** command will hang forever. You must apply an AIX LVM PTF to correct this problem before you can implement the extendvg function. See "Host system requirements" on page 10 for more information about applying PTFs.

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 67.

Creating enhanced concurrent-capable volume groups

Perform 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>lvlstmajor
41,54..58,60..62,67,78...
monkey>lvlstmajor
39..49,55,58,67,80...
zebra>lvlstmajor
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 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 the entry fields.
Press Enter AFTER making all desired changes.

 [Entry Fields]
VOLUME GROUP name                [con_vg]
Physical partition SIZE in megabytes            4
Activate volume group AUTOMATICALLY at system restart?     no
Volume Group MAJOR NUMBER             [80]
Create VOLUME GROUPS concurrent-capable?         yes
Auto-varyon in concurrent mode?              no
LTG size in kbytes                  128
********************************************************************************
```

Importing enhanced concurrent-capable volume groups

Perform the following step to import enhanced concurrent-capable volume groups.

Enter **smitty importvg** at the command prompt.

Notes:

1. Before importing enhanced concurrent-capable volume groups on SDD devices, issue the **lspv** command to make sure there is pvid on the vpath. If pvid is not

displayed, import the volume group on one of the hdisks that belongs to the vpath. Enter **hd2vp** to convert the volume group to vpaths.

2. If the hdisks do not have a pvid, issue the **chdev -l hdiskX -a pv=yes** to recover it. To verify that pvid now exists, issue the **lspv** command against the hdisk. This method can also be used when attempting to obtain a pvid on a vpath.
3. Verify that the volume group is not varied on for any of the nodes in the cluster prior to attempting to retrieve the pvid.
4. A screen similar to the following 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 an SDD *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 vpath/hdisks on all nodes in the cluster.

Perform 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 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]
********************************************************************************
```

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

Perform 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 an SDD *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 fallover

Typically, when there is a node failure, HACMP transfers ownership of shared disks and other resources through a process known as node fallover. Certain situations, such as a loose or disconnected SCSI or fibre-channel-adapter card, can cause your vpath devices to lose one or more underlying paths during node fallover. Perform the following steps to recover these paths:

- Check to ensure that all the underlying paths (hdisks) are in the *Available* state.
- Enter the **addpaths** command to add the lost paths back to the SDD devices.

If your 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 devices of a volume group” on page 41 for more information about the **addpaths** command.

Note: Running the **cfgmgr** command while the 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 64.

Managing secondary-system paging space

SDD 1.3.2.6 (or later) supports secondary system paging on *multipath* fibre-channel vpath device from an AIX 4.3.3 or AIX 5.1.0 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.

Attention: 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 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 -s*NN 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.
-s*NN*	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:

rmps *PS01*

Providing load-balancing and failover protection

SDD provides load-balancing and failover protection for AIX applications and for the LVM when supported storage device 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 vpath device configuration

To provide failover protection, a supported storage device must have a minimum of two paths. Both the SDD vpath device and the supported storage device 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 vpath devices are available to provide failover protection, use either the Display Data Path Device Configuration SMIT panel, or run the **lsvpcfg** command. Perform the following steps to use 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 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 pseudo 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:

```
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 or not 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 86 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 291 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 69) and then add paths later by reconfiguring the supported storage device. (See “Dynamically adding paths to SDD devices of a volume group” on page 41 for information about adding paths to a SDD device.) However, an SDD volume group does not have failover protection if any of its physical volumes has only a single path.

Perform the following steps to create a new volume group with SDD vpaths:

1. Enter **smitty** at the AIX command prompt. The System Management Interface Tool (SMIT) is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.
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 vpaths.

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 69, 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:

```
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 disk (or pseudo) 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:

- Through the loss of a device path
- By creating a volume group from single-path vpath (pseudo) devices
- As a side effect of running the disk change method
- Through running the **mksysb restore** command
- By manually deleting devices and running the configuration manager (cfgmgr)

The following sections provide more information about the ways that failover protection can be lost.

Through the loss of a device path

Due to hardware errors, SDD might remove one or more nonfunctional paths from a vpath pseudo device. The states of these nonfunctional paths are marked as Dead, Invalid, or Close_Dead by SDD. A pseudo device will lose failover protection if it has only one functional path left. To determine if any of the pseudo devices have lost failover protection due to nonfunctional paths, use the **datapath query device** command to show the state of paths to a pseudo device.

By creating a volume group from single-path vpath (pseudo) devices

A volume group created using any single-path pseudo device does not have failover protection because there is no alternate path to the supported storage device LUN.

As 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 vpath pseudo device and one or more of its hdisks to have the same pvid attribute in the ODM. If the volume group containing the vpath pseudo 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 queue_depth=30
```

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 pv 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 vpath pseudo 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 **dpovgfix vg-name** command after running the **chdev** command.

By manually deleting devices and running the configuration manager (cfgmgr)

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, only one path of the original vpath3 is configured by AIX. The following commands would produce this situation:

```
rmdev -dl vpath3
rmdev -dl hdisk4
rmdev -dl hdisk27
cfgmgr
```

The **datapath query device** command displays the vpath3 configuration state.

Next, all paths to the vpath must be restored. You can restore the paths in one of the following ways:

- Enter **cfgmgr** once for each installed SCSI or fibre-channel adapter.
- Enter **cfgmgr** *n* times, where *n* represents the number of paths per SDD device.

  Tip: Running the AIX configuration manager (**cfgmgr**) *n* times for n-path configurations of supported storage devices is not always required. It depends on whether the supported storage device has been used as a physical volume of a volume group or not. If it has, it is necessary to run **cfgmgr** *n* times for an n-path configuration. Because the supported storage device has been used as a physical volume of a volume group before, it has a pvid value written on its sector 0.

  When the first SCSI or fibre-channel adapter is configured by **cfgmgr**, the AIX disk driver configuration method creates a pvid attribute in the AIX ODM database with the pvid value it read from the device. It then creates a logical name (hdisk*N*), and puts the hdisk*N* in the *Defined* state. When the second adapter is configured, the AIX disk driver configuration method reads the pvid from the same device again and searches the ODM database to see if there is already a device with the same pvid in the ODM. If there is a match, and that hdisk*N* is in a `Defined` state, the AIX disk driver configuration method does not create another hdisk logical name for the same device. That is why only one set of hdisks gets configured the first time **cfgmgr** runs. When **cfgmgr** runs for the second time, the first set of hdisks are in the `Available` state, so a new set of hdisks are *Defined* and configured to the *Available* state. That is why you must run **cfgmgr** *n* times to get *n* paths configured. If the supported storage device has never belonged to a volume group, that means there is no pvid written on its sector 0. In that case, you need to run the **cfgmgr** command only once to get all multiple paths configured.

  Note:

  The **addpaths** command allows you to dynamically add more paths to SDD devices while they are in *Available* state. The **cfgmgr** command must always be run *N* times when adding new LUNs.

  In addition, this command allows you to add paths to vpath devices (which are then opened) belonging to active volume groups.

  This command will open a new path (or multiple paths) automatically if the vpath 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 devices of a volume group” on page 41 for more information about the **addpaths** command.

The following command shows an example of how to unconfigure an SDD device to the *Defined* state using the command-line interface:

```
rmdev -l vpathN
```

The following command shows an example of how to unconfigure *all* SDD devices to the *Defined* state using the command-line interface:

```
rmdev -l dpo -R
```

The following command shows an example of how to configure a vpath device to the *Available* state using the command-line interface:

```
mkdev -l vpathN
```

The following command shows an example of how to configure all vpath devices to the *Available* state using the command-line interface: state using the SMIT:

```
smitty device
```

The following command shows an example of how to configure all vpath devices to the *Available* state using the command-line interface:

```
cfallvpath
```

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.

Perform the following steps to import a volume group with SDD devices:

1. Enter **smitty** from your desktop window. SMIT is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.
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, perform the following tasks:
   - Enter the volume group that you want to import.
   - Enter the physical volume that you want to import.
   - Press **Enter** after making the changes.

   You can press **F4** for a list of choices.

Importing volume groups under HACMP

Attention: SDD does not automatically create the pvid attribute in the ODM database for each vpath device. The AIX disk driver automatically creates the pvid attribute in the ODM database, if a pvid exists on the physical device. Therefore, the first time 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 85) to convert the volume group's physical volumes from supported storage device hdisks to vpath devices. This conversion step not only creates pvid attributes for all vpath devices that belong to that imported volume group, it also deletes the pvid attributes for these vpath devices' underlying hdisks. Later on, you can import and vary on the volume group directly from the vpath devices. These special requirements apply to both concurrent and nonconcurrent volume groups.

See “Special requirements” on page 58 for special requirements regarding unconfiguring and removing SDD vpath devices when running in HACMP environments.

See page 59 for an example of hdisks' pvids that do not exist.

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.

Perform the following steps to export a volume group with SDD devices:

1. Enter **smitty** from your desktop window. SMIT is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.

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

Run the **dpovgfix** shell script to recover a mixed volume group. The syntax is **dpovgfix vg-name**. The script searches for a pseudo device corresponding to each hdisk in the volume group and replaces the hdisk with the vpath pseudo device. In order for the shell script to be executed, 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 69) 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 devices of a volume group” on page 41 for information about adding paths to an SDD device.)

Perform the following steps to extend a volume group with SDD devices:

1. Enter **smitty** from your desktop window. SMIT is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.
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 83.

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 devices, go to "Accessing the Remake a Volume Group with Data Path Devices SMIT panel" on page 84.

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 19 lists the SDD-specific SMIT panels and how you can use them.

Table 19. SDD-specific SMIT panels and how to proceed

SMIT panels	**How to proceed** **Go to:**
Display Data Path Device Configuration	"Accessing the Display Data Path Device Configuration SMIT panel" on page 79
Display Data Path Device Status	"Accessing the Display Data Path Device Status SMIT panel" on page 80
Display Data Path Device Adapter Status	"Accessing the Display Data Path Device Adapter Status SMIT panel" on page 80
Define and Configure all Data Path Devices	"Accessing the Define and Configure All Data Path Devices SMIT panel" on page 81
Add Paths to Available Data Path Devices	"Accessing the Add Paths to Available Data Path Devices SMIT panel" on page 81

Table 19. SDD-specific SMIT panels and how to proceed (continued)

Configure a Defined Data Path Device	"Accessing the Configure a Defined Data Path Device SMIT panel" on page 81
Remove a Data Path Device	"Accessing the Remove a Data Path Device SMIT panel" on page 82
Add a Volume Group with Data Path Devices	"Accessing the Add a Volume Group with Data Path Devices SMIT panel" on page 82
Add a Data Path Volume to a Volume Group	"Accessing the Add a Data Path Volume to a Volume Group SMIT panel" on page 82
Remove a Physical Volume from a Volume Group	"Accessing the Remove a Physical Volume from a Volume Group SMIT panel" on page 83
Back Up a Volume Group with Data Path Devices	"Accessing the backup a Volume Group with Data Path Devices SMIT panel" on page 83
Remake a Volume Group with Data Path Devices	"Accessing the Remake a Volume Group with Data Path Devices SMIT panel" on page 84

Accessing the Display Data Path Device Configuration SMIT panel

Perform the following steps to access the Display Data Path Device Configuration panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.
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.                       |
|                                                                    |
|                                       [Entry Fields]               |
| Select Query Option                    all devices           +     |
| Device Name/ Device Model                    [ ]                   |
+--------------------------------------------------------------------+
```

The Select Query Option has three options:

All devices
This option executes **lsvpcfg** and all the data path devices are displayed. No entry is required in the Device Name/Device Model field.

Device name
This option executes **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 executes **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 86 for detailed information about the **lsvpcfg** command.

Accessing the Display Data Path Device Status SMIT panel

Perform the following steps to access the Display Data Path Device Status panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.
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.                       |
|                                                                    |
|                                        [Entry Fields]              |
| Select Query Option                     all devices            +   |
| Device Number/ Device Model                   [ ]                  |
+--------------------------------------------------------------------+
```

The Select Query Option has 3 options:

All devices
This option executes **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 executes **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 executes **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 291 for detailed information about the **datapath query device** command.

Accessing the Display Data Path Device Adapter Status SMIT panel

Perform the following steps to access the Display Data Path Device Adapter Status panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
- The list items on the SMIT panel might be worded differently from one AIX version to another.

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

To access the Define and Configure All Data Path Devices panel, perform the following steps:

1. Enter **smitty** from your desktop window. SMIT is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.

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

Perform the following steps to access the Add Paths to Available Data Path Devices panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.

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

Perform the following steps to access the Configure a Defined Data Path Device panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.

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

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

Perform the following steps to access the Remove a Data Path Device panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.

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

Perform 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.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.

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

Perform 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.

   Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
- The list items on the SMIT panel might be worded differently from one AIX version to another.

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

Perform 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.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.

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

Perform 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.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.

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, perform the following steps:
   - Enter the Backup DEVICE or FILE name.
   - Enter the Volume Group to backup.
   - 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

Perform 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.

   Tip:
   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.
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

addpaths

You can use the **addpaths** command to dynamically add more paths to SDD devices when they are in the *Available* state. In addition, this command allows you to add paths to vpath devices (which are then opened) belonging to active volume groups.

This command will open a new path (or multiple paths) automatically if the vpath is in *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.

The syntax for this command is:

For more information about this command, go to “Dynamically adding paths to SDD devices of a volume group” on page 41.

hd2vp and vp2hd

SDD provides two conversion scripts, hd2vp and vp2hd. The hd2vp script converts a volume group from supported storage device hdisks to SDD vpaths, and the vp2hd script converts a volume group from SDD vpaths 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 SDD from your AIX host system.

Note: You must convert all your applications and volume groups to the original supported storage device hdisk device special files before removing SDD.

The syntax for these conversion scripts is as follows:

vgname
Specifies the volume group name to be converted.

These two conversion programs require that a volume group contain either *all* original supported storage device hdisks or *all* SDD vpaths. The program fails if a volume group contains both kinds of device special files (mixed volume group).

Tip: Always use SMIT to create a volume group of SDD devices. This avoids the problem of a mixed volume group.

dpovgfix

You can use the dpovgfix script tool to recover mixed volume groups.

Performing AIX system management operations on adapters and supported storage device hdisk devices can cause original supported storage device hdisks to be contained within a 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 hdisk*N* -a queue_depth=30

If this disk is an active hdisk of a vpath 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 vpath device. The result is that an SDD volume group partially uses 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 command syntax is:

vgname
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 command syntax is:

   lsvpcfg

   See "Verifying the SDD configuration" on page 38 for an example of the output and what it means.
2. The command can also be issued using the vpath device name as a parameter. The command syntax is:

   lsvpcfg vpath*N*$_0$ vpath*N*$_1$ vpath*N*$_2$

   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 38 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 a vpath. The command syntax is:

   lsvpcfg device model

   The following are examples of valid device models:

   2105 Display all 2105 models (ESS).

   2105F Display all 2105 F models (ESS).

   2105800
   Display all 2105 800 models (ESS).

2145 Display all 2145 models (SAN Volume Controller).

2062 Display all 2062 models (SAN Volume Controller for Cisco MDS 9000).

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 70.

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 77.

querysn

You can use the **querysn** command to get the serial number for the logical device (hdisk) and save it into an exclude file (/etc/vpexclude). During the SDD configuration, the SDD configure methods read all serial numbers listed in this file and exclude these devices from the SDD configuration.

The syntax is:

```
►►—querysn <-d> -l— device-name—————————————————————————►◄
```

-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 **querysn** command deletes all existing contents from this exclude file and then writes the new serial number into the file; otherwise, it appends the new serial number at the end of the file.

device name
Specifies the SDD device name of the supported storage device (hdiskN).

Example:

```
querysn -l hdisk10
```

Notes:

1. Do not use the **querysn** 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 **querysn** command on only one of the logical names of that LUN.
3. You should not use the **querysn** command on the same logical device multiple times. Using the **querysn** command on the same logical device multiple times results in duplicate entries in the /etc/vpexclude file, and the system administrator will have to administer the file and its content.
4. Executing the **querysn** command with the -d flag deletes all existing contents from the exclude file and then writes the new serial number into the file. If you want to remove only one device from the /etc/vpexclude file, you must edit the /etc/vpexclude with the vi editor and delete the line containing the device name. To replace a manually excluded device in the SDD configuration, you have to open the /etc/vpexclude file with a text editor (for example, vi) and delete the

line containing the device name. For detailed instructions on the proper procedure, see “Replacing manually excluded devices in the SDD configuration” on page 19.

Persistent reserve command tool

With SDD 1.3.2.9 (or later), the **lquerypr** 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 registrations

With SDD 1.4.0.0 or later, this command can be issued to both SDD vpath devices and supported storage device hdisk devices.

It can be used as a tool for the user in the situation that SDD persistent reserve version is installed, however, HACMP is not installed on multiple AIX servers, or on a server with multiple logical partitions (LPAR) configured and sharing disk resources in nonconcurrent mode.

In the case that the primary resource owner suddenly goes down without releasing the persistent reserve, and for some reason it cannot be brought up for a while, then the standby node of LPAR or server cannot take the ownership of sharing resources.

Notes:

1. 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.
2. If you are running in a SAN File System environment, there might be special restrictions and considerations regarding use of SCSI Persistent Reserve or SCSI Reserve. Please consult the SAN File System documentation shown in “The SAN File System library” on page xxiii for more information.

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

lquerypr [[-p]|[-c]|[-r]][-v][-V][-h/dev/PVname]

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 or read

reservation key, release persistent reservation, preempt-abort persistent reservation, and clear persistent reservation.

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 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, then 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

After you configure the SDD, it creates SDD devices (vpath devices) for supported storage device LUNs. Supported storage device LUNs are accessible through the connection between the AIX host server SCSI or FCP adapter and the supported storage device ports. The AIX disk driver creates the original or supported storage devices (hdisks). Therefore, with SDD, an application now has two ways in which to access supported storage devices.

To use the SDD load-balancing and failover features and access supported storage devices, your application must use the SDD vpath devices rather than the supported storage device hdisk devices.

Two types of applications use supported storage device disk storage. One type of application accesses supported storage devices through the SDD vpath device (raw device). The other type of application accesses supported storage devices through the AIX Logical Volume Manager (LVM). For this type of application, you must create a volume group with the SDD vpath devices.

Note: If you are running in a SAN File System environment, there might be special restrictions and considerations regarding use of any Logical Volume Manager such as LVM. Please consult the SAN File System documentation shown in "The SAN File System library" on page xxiii for more information.

If your application used supported storage device hdisk device special files directly before installing SDD, convert the application to use SDD vpath device special files. After installing SDD, perform the following steps:

1. Enter **smitty** from your desktop window. SMIT is displayed.

   Tip:

   - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
   - The list items on the SMIT panel might be worded differently from one AIX version to another.
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 vpaths with their attached multiple paths (hdisks), select all 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 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

Attention: You must use SMIT. The SMIT facility runs in two interfaces, nongraphical (enter **smitty** to invoke the nongraphical user interface) or graphical (enter **smit** to invoke the GUI).

Do not use the **mkvg** command directly; the path failover capability could be lost.

If your application accesses supported storage devices through LVM, determine the volume group that the application uses before you convert volume groups. Then, perform the following steps to convert the volume group from the original supported storage device hdisks to the SDD vpaths:

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, perform 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 vpaths:

   ```
   hd2vp vgname
   ```
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 22.
- The supported storage subsystem 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.

    Tip:
    - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
    - The list items on the SMIT panel might be worded differently from one AIX version to another.
  - 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 pseudo 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 vpath 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  hdiskN  -a  pv=clear
  chdev -l  vpathN  -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.

      Tip:
      - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
      - The list items on the SMIT panel might be worded differently from one AIX version to another.

   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 devices from the volume group. The Remove a Physical Volume panel is displayed. Remove all non-SDD devices.

Note: A non-SDD volume group can consist of non-ESS, non-SAN Volume Controller non-SAN Volume Controller for Cisco MDS 9000, ESS, or SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 hdisk devices. There is no failover protection unless multiple paths are configured for each LUN.

Migrating an existing non-SDD volume group to SDD vpath devices in concurrent mode

This procedure shows how to migrate an existing AIX volume group to use SDD vpath (pseudo) 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 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, we have a pseudo device, vpath12, with two paths, hdisk14 and hdisk30, that we will migrate the volume group to.

1. Add the 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.

      Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can enter **smit** to invoke the GUI.
- The list items on the SMIT panel might be worded differently from one AIX version to another.

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 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:

```
rmlvcopy loglv01 1 hdisk13
rmlvcopy 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.

If you are running SDD 1.4.0.0 (or later):

1. Enter **pathtest -d** . (for example, **pathtest -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.

Error messages

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 following list shows the SDD error messages and explains each one:

VPATH_XBUF_NOMEM
An attempt was made to open an SDD vpath 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 device file failed to open one of its paths (hdisks). An attempt to open a vpath device is successful if at least one attached path opens. The attempt to open a vpath device fails only when *all* the vpath device paths fail to open.

VPATH_DEVICE_OFFLINE
Several attempts to retry an I/O request for a 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 12, “Using the datapath commands,” on page 283.

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
A vpath device has no path available for an I/O operation. The state of the vpath is set to LIMBO. All following I/Os to this vpath device are immediately returned to the caller.

Error messages for the HACMP environment

The following list shows the error messages logged by SDD in a HACMP environment. See “SDD persistent reserve attributes” on page 58 for more information about persistent reserve.

VPATH_FAIL_RELPRESERVE
An attempt was made to close a vpath device that was not opened with the **RETAIN_RESERVE** option on the persistent reserve. The attempt to close the 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 a vpath device, but the reservation key of the 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.

Error messages for AIX Hot Plug support

The following error messages are available with SDD 1.5.1.0 or later in an AIX Hot Plug (AIX 5L or later) supported environment:

VPATH_ADPT_SUSPENDED
The **datapath set adapter** *n* **suspend** command was executed. The suspended adapter and its child devices in SDD are not operational for I/O.

VPATH_ADPT_RESUMED
The **datapath set adapter** *n* **resume** command was executed. The suspended adapter and its child devices in SDD are now operational for I/O.

VPATH_ADPT_REMOVED
The **datapath remove adapter** *n* command was executed. Adapter *n* and its child devices are removed from SDD.

VPATH_PATH_REMOVED
The **datapath remove device** *m* **path** *n* command was executed. Path *n* for device *m* is removed from SDD.

Chapter 3. Using SDDPCM on an AIX host system

SDDPCM is a loadable path control module for ESS devices to supply path management functions and error recovery algorithms. When the ESS devices are configured as MPIO-capable devices, SDDPCM is loaded during the ESS device configuration and becomes part of the AIX MPIO SCSI/FCP (Fibre Channel Protocol) device driver. The AIX MPIO-capable device driver with the ESS SDDPCM module enhances the data availability and I/O load balancing.

This chapter provides step-by-step procedures on how to install SDDPCM on an AIX 5.2.0 ML3 (or later) host system, and configure ESS devices into Multipath I/O (MPIO)-capable devices.

For detailed information about AIX 5.2.0 ML3 (or later) MPIO support, visit the following Web site:

http://publib16.boulder.ibm.com/pseries/en_US/aixbman/baseadmn/manage_MPIO.htm

AIX MPIO-capable device drivers will 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 ESS licensed internal code
- Prevention of a single failing bus adapter on the host system, SCSI, fibre channel cable, or host-interface adapter on the ESS from disrupting data access

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

www-1.ibm.com/servers/storage/support/software/sdd.html

SDD and SDDPCM are exclusive software packages on a server for one kind of storage. You cannot install both software packages on a server for ESS devices. When ESS 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 should install SDDPCM in order to configure ESS devices into MPIO-capable-devices (where only one logical device instance is created for a physical LUN). In order to run SDDPCM on AIX 5.2.0 ML3 and later, you must install all the latest PTFs for that OS level.

Configuring ESS devices into MPIO-capable devices or into non-MPIO-capable devices is controlled by ESS host attachment.

To configure ESS devices as non-MPIO-capable devices, install the ibm2105.rte package with a version of 32.6.100.x

To configure ESS devices as MPIO-capable devices, install the devices.fcp.disk.ibm2105.mpio.rte package with a version of 33.6.100.y

For the latest version of the ESS host attachment package, refer to the Readme file on the SDD download web site:

© Copyright IBM Corp. 1999, 2004

www-1.ibm.com/servers/storage/support/software/sdd.html

Note: ESS SCSI devices are not supported by SDDPCM, and the ESS host attachment for SCSI MPIO support (devices.scsi.disk.ibm2105.mpio.rte) is not being shipped.

If a server connects to both ESS and SAN Volume Controller, then you can install SDDPCM for ESS, and install SDD for SAN Volume Controller to support coexistence

Supported SDDPCM features

The following SDDPCM features are supported in this release:

- 32- and 64-bit kernels
- Automatic failed paths reclamation by healthchecker
- Four types of reserve policies are supported:
  - No_reserve policy
  - Exclusive host access single path policy
  - Persistent reserve exclusive host policy
  - Persistent reserve shared host access policy
- Three path-selection algorithms are supported:
  - Failover
  - Round robin
  - Load balancing
- Failback error-recovery algorithm
- Fibre-channel dynamic device tracking
- Support for the ESS Model 750 device module
- Support for an external ESS MPIO device as the system boot device
- Support for an external ESS MPIO device as the primary or secondary dump device
- ESS multipath devices as system paging space
- Support for the SDDPCM server daemon for the enhanced path health check function
- Support for Web-based System Manager (WebSM) for ESS MPIO devices. (Refer to www-1.ibm.com/servers/aix/wsm/ for more information about WebSM).
- Support for a maximum of 1200 LUNs
- Dynamically adding paths or adapters
- Dynamically removing paths or adapters
- **iostat** command with new command options in AIX 5.2.0 ML3 (or later)
- Last path of device is reserved and is never placed into the Failed state
- Support for the essutil Product Engineering tool into SDDPCM's pcmpath program

Unsupported SDDPCM features

The following SDDPCM features are not currently supported. Support for these features will be added in future releases.

- HACMP
- GPFS

- Support for SAN Volume Controller
- Support for SAN Volume Controller for Cisco MDS 9000

Verifying the hardware and software requirements

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

Hardware

- ESS
- One or more switches, if the ESS is not direct-attached
- Host system
- Fibre-channel adapters and cables

Software

- AIX 5.2.0 ML3 (or later) operating system
- Fibre-channel device drivers
- devices.sddpcm.52b.rte (version 2.0.1.0) installation package
- ESS host attachment for SDDPCM (devices.fcp.disk.ibm2105.mpio.rte).

Unsupported environments

SDDPCM does not support:

- A host system with both a SCSI and fibre-channel connection to a shared ESS logical unit number (LUN)
- Single-path mode during concurrent download and installation of LIC nor during any ESS concurrent maintenance that impacts the path attachment, such as an ESS host-bay-adapter replacement.

Host system requirements

To successfully install SDDPCM for ESS, you must have AIX 5.2.0 ML3 (or later) installed on your host system along with the AIX required fixes, APARs, and microcode updates identified on the following Web site:

www-1.ibm.com/servers/storage/support/

ESS requirements

To successfully install SDDPCM, ensure that the correct ESS attachment package is installed on the server:

- devices.fcp.disk.ibm2105.mpio.rte (version 33.6.100.y)

Fibre requirements

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

www-1.ibm.com/servers/eserver/support/

If your host has only one fibre-channel adapter, it requires you to connect through a switch to multiple ESS 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 Web site:

www-1.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 pSeries with AIX 5.2.0 ML3 (or later).
- 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 an ESS 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 ESS 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 ESS 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 health-check function
sample_pcmsrv.conf	The sample SDDPCM server daemon configuration file
fcppcmmap	Collects ESS fibre-channel device information through SCSI commands

Before you install SDDPCM, you must perform the tasks identified in the following section:

Preparing for SDDPCM installation on an ESS

Before you install SDDPCM, you must configure:

- The ESS to your host system and the required fibre-channel adapters that are attached.
- The ESS for single- or multiple-port access for each LUN. SDDPCM requires a minimum of two independent paths that share the same logical unit to use the load-balancing and failover features. **With a single path, failover protection is not provided.**

For more information about how to configure your IBM ESS, refer to the *IBM TotalStorage Enterprise Storage Server: Introduction and Planning Guide*.

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), if it is installed.
- Install the AIX fibre-channel device drivers, if necessary.
- Verify and upgrade the fibre channel adapter firmware level
- Install the MPIO-supported ESS attachment (version 33.6.100.y)

Determining the correct installation package

SDDPCM can be installed only on an AIX 5.2.0 ML3 (or later) operating system. The package name of SDDPCM is devices.sddpcm.52b.rte.

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 ESS device configuration, then you must unconfigure and remove all SDD vpath devices, and then uninstall the SDD package. See "Unconfiguring SDD" on page 31 and "Removing SDD from an AIX host system" on page 39.

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, then you must:
   a. Unconfigure and remove all ESS hdisk devices.
   b. Use **smitty** to uninstall the ibm2105.rte package.

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 Web site:

www-1.ibm.com/servers/storage/support/

Perform 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 pSeries with feature code 6239.

devices.pci.df1000f9
The adapter device driver for RS/6000 or pSeries with feature code 6228.

devices.pci.df1000f7
The adapter device driver for RS/6000 or pSeries 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 -i | grep IY*nnnnn*

    where ***nnnnn*** represents the APAR numbers.

    If the APARs are listed, that means that they are installed. If they are installed, go to “Configuring ESS MPIO-capable devices” on page 108. 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.

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 Web site:

https://techsupport.services.ibm.com/server/mdownload

Tip:

- The current firmware level for LP7000E adapter is sf322A1.
- The current firmware level for LP9002 adapter is sf382A1.

Perform 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).............S2F3.22A1
```

   To verify the firmware level, ignore the second character in the **ZB** field. In the example, the firmware level is sf322A1.
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 adapter firmware level."

Upgrading the 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 15). 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.

Installing the MPIO-supported ESS attachment

You must install the MPIO-supported ESS attachment before devices.sddpcm.52b.rte is installed. Otherwise, the SDDPCM installation will fail. The attachment VRMF starts from 33.6.100.y.

The devices.fcp.disk.ibm2105.mpio.rte device-attachment package for ESS fibre-channel devices is provided. This package must be installed before you install the device.sddpcm.52b.rte package.

Installing SDDPCM

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.

Notes:

1. To mount the CD-ROM and install SDDPCM, you must have root access and AIX system administrator knowledge.
2. The devices.fcp.disk.ibm2105.mpio.rte (for ESS FCP devices) package must be installed before you install the device.sddpcm.52b.rte package.

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 perform 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.

1. Log in as the root user
2. Insert the compact disc into the CD-ROM drive.
3. From your desktop window, enter **smitty fs** and press **Enter**.
4. Select **Add / Change / Show / Delete File Systems** and press **Enter**.
5. Select **CDROM File System** and press **Enter**.
6. Select **Add a CDROM File System** and press **Enter**. The Add a CDROM File System panel is displayed.
7. Select **DEVICE name** and select **F4**. The DEVICE name panel is displayed.
8. Select the compact disc drive that you are using for the installation, (for example, cd0), and press **Enter**.
9. Select **MOUNT POINT** and enter a directory where you want the CDROM File System to be mounted, (for example, /cdmnt).
10. Click the default option settings for the other fields to ensure that they are want you need.

```
+-------------------------------------------------------------+
+ Add a CDROM File System                                     +
+                                                             +
+ Type or select values in entry fields.                     +
+ Press Enter AFTER making all desired changes.              +
+                                                             +
+                                          [Entry Fields]  +
+ * DEVICE name                           cd0             +
+ * MOUNT POINT                           [/cdmnt]        +
+ Mount AUTOMATICALLY at system restart?     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.

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. Perform the following SMIT steps to install the SDDPCM package on your system.

1. From your desktop window, cd to the directory where the CDROM 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.52b.rte installation package.
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 umount 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.

Verifying the currently installed version of SDDPCM

You can verify your currently-installed version of SDDPCM by issuing the following command:

lslpp -l *sddpcm*

Maximum number of devices supported by SDDPCM

SDDPCM supports a maximum of 1200 configured devices and a maximum of 16 paths per device. However, with the round robin or load balance path selection algorithms, configuring more than four paths per device may impact the I/O performance. You should use the minimum number of paths necessary to achieve sufficient redundancy in the SAN environment. The recommended number of path per device is two or four.

In order to support 1200 ESS LUNs, system administrators should first determine whether the system has sufficient resources to support a large number of devices. See “Determining whether system has enough resource for increasing the maximum number of ESS LUNs” on page 26 for more information. The AIX disk driver APAR IY49825 must be installed on AIX 5.2.0 ML3 or later.

Configuring and unconfiguring ESS MPIO-capable devices

After installing MPIO-supported ESS host attachment and the SDDPCM package, you need to reboot the system in order to configure ESS device as MPIO-capable devices. After the first system reboot, you can then use the normal AIX command line configure programs to configure and unconfigure ESS MPIO-capable devices.

After the system reboots, the SDDPCM server daemon (pcmsrv) should automatically start.

Configuring ESS MPIO-capable devices

The newly installed ESS 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, the **cfgmgr** command must be executed once for each host adapter each time a device is added.

If you use the **cfgmgr** command to configure ESS MPIO devices, you might need to start the SDDPCM server daemon manually, if it has not already started. See "SDDPCM server daemon" on page 136 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 reboots, the SDDPCM server daemon (pcmsrv) should automatically start.

Unconfiguring ESS MPIO-capable devices

To remove all ESS MPIO-capable devices:

1. Unmount the file systems of all ESS devices.
2. Vary off all ESS device volume groups.
3. Enter the **stopsrc -s pcmsrv** command to stop pcmsrv.
4. Enter the following command for each adapter:

   rmdev -dl fcsX -R

   Note: This command requires that PTF U488799 is installed.

Verifying the SDDPCM Configuration

To verify the SDDPCM configuration, you can use one of the following:

- SMIT MPIO management submenu, or
- SDDPCM **pcmpath query device** command

Perform the following steps use SMIT to verify the SDDPCM configuration on an AIX host system:

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

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 FC 2105XXX MPIO devices and ensure that they are in the Available state.

You can also use the SDDPCM **pcmpath query device** command to check the configuration status of ESS devices.

Note: If none of the ESS devices are configured successfully as MPIO devices, then the **pcmpath query device** command will fail.

Updating SDDPCM

The following sections discuss the following methods of updating SDDPCM:

- "Updating SDDPCM packages by applying a program temporary fix" on page 110
- "Committing or rejecting a program temporary fix update" on page 111
- "Migrating from SDDPCM to the AIX default PCM or to SDD" on page 112

Updating SDDPCM packages by applying a program temporary fix

SDDPCM allows you to update SDDPCM by installing a PTF. A PTF file has a file extension of .bff (for example, devices.sddpcm.52b.rte.2.0.1.1.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.

Before applying the PTF to your system, you must unconfigure all ESS devices from the *Available* state to the *Defined* state and you must stop the SDDPCM server daemon. After applying the PTF, follow the procedure in "Configuring and unconfiguring ESS MPIO-capable devices" on page 108 to reconfigure the ESS devices. You must also restart the SDDPCM server daemon.

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).

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

If the PTF is on a CD, you need to mount the CD file system, and then 'cd' to the directory on the CD that contains the SDDPCM PTF. See "Creating and mounting the CD-ROM filesystem" on page 106 for directions on how to mount the CD file system. Throughout this SMIT procedure, /dev/cd0 is used for the CD drive address. The drive address can be different in your environment.

Perform the following SMIT steps to update the SDDPCM package on your system:

1. Log in as the root user.
2. 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.
3. Select **Install Software** and press **Enter**.
4. Enter **.** to select the current directory as the INPUT Device/Directory for Software panel and press **Enter**. The Install Software panel is displayed.
5. Select **Software to Install** and press **F4**. The Software to Install panel is displayed.
6. Select the PTF package that you want to install.
7. Press **Enter**. The Install and Update from LATEST Available Software panel is displayed with the name of the software that you selected to install.
8. 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?.
9. Check the other 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.
13. Unmount the CD file system and remove the compact disc.

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

Committing or rejecting a program temporary fix update

Before you reject a PTF update, you need to unconfigure and remove all ESS devices from your host system. Committing a PTF does not require this extra step. Perform 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).

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

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 the system even though the bosboot message may indicate that a reboot is necessary.

Configuring ESS MPIO-capable devices as the boot device

An ESS MPIO-capable device can be used as the system boot device. To configure the ESS boot device with the SDDPCM module:

1. Select an ESS device as the boot device.
2. Install AIX 5.2.F (or later) operating system on the selected ESS device.
3. Reboot the system. The ESS boot device is configured as an MPIO-capable device with AIX default PCM.
4. Install the ESS host attachment for SDDPCM and SDDPCM packages.
5. Reboot the system.

All ESS MPIO-capable devices, including ESS MPIO boot devices, are now configured with SDDPCM.

Migrating from SDDPCM to the AIX default PCM or to SDD

Note: If you have ESS MPIO boot devices configured with SDDPCM, then migration from SDDPCM to the AIX default PCM is not supported in this release.

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 116 for directions on uninstalling SDDPCM. After you uninstall SDDPCM, you can then reboot the system to migrate ESS MPIO devices to the AIX default PCM. If you want to migrate ESS devices to SDD devices, you must then install the ESS host attachment for SDD (ibm2105.rte) and the appropriate SDD package for your system. Then reboot the system to configure the ESS devices to SDD vpath devices.

Support system dump device with the ESS MPIO-capable device

You can choose an ESS MPIO-capable device to configure with the system primary and secondary dump devices. You can configure the system dump device with the ESS boot device, or with the non-boot device. The path selection algorithm for the system dump device will automatically default 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, then I/O is routed to the next path being selected.

SDDPCM ODM attribute settings

The following sections discuss the SDDPCM ODM attribute default settings, and how to change the attributes of the ESS MPIO-capable devices:

- “SDDPCM ODM attribute default settings”
- “Changing device reserve policies”
- “Changing the path selection algorithm” on page 113
- “Changing SDDPCM path healthcheck mode” on page 113

SDDPCM ODM attribute default settings

SDDPCM has following default attributes settings:

Attribute	Default value
device reserve policy	no_reserve
path selection algorithm	load balance
healthcheck mode	nonactive
healthcheck time interval	20 seconds

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 100):

- 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 path selection algorithm supported 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.

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, then 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: The current AIX operating system does not provide a way for you to generate, query, or remove a persistent reserve key on a LUN. If you want to switch to the persistent reserve type of reserve policy, you must provide a persistent reserve key. Currently, persistent reserve policy is not a recommended reserve policy.

Changing the path selection algorithm

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

Changing SDDPCM path healthcheck mode

SDDPCM supports the path healthcheck function. If this function is enabled, then SDDPCM will test opened paths and reclaim failed paths based on the value set in the following device healthcheck attribute:

hcheck_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 will be 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, then the healthcheck command will only be sent to failed paths, because the round robin and load balanced algorithms route I/O to all opened paths that are functional. The default value setting of SDDPCM is *nonactive*.

Use the following command to change the device path healthcheck mode to failed:

chdev -l hdiskX -a hcheck_mode=failed

Changing SDDPCM path healthcheck time interval

The hcheck_interval attribute will determine how often the paths of a device should be health-checked. The hcheck_interval attribute has a range of values from 0 - 3600 seconds. When a value of 0 is selected, the healthcheck function is disabled. The default value setting is 20 (seconds).

Use the following command to disable the device healthcheck function:

chdev -l hdiskX -a hcheck_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 the SDDPCM healthcheck function is disabled, then the SDDPCM server daemon will also healthcheck failed paths that are already opened. See Chapter 11, “Using the SDD server and the SDDPCM server,” on page 279 for more information.

Dynamically enabling and disabling paths or 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 **pcmpath** command.

   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 5.2.0 ML3 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:

   chpath -l hdiskX -s E|D -p fscsiX -w
   ″5005076300c99b0a,5200000000000000″
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.

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 adaper N offline

Note: 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

Dynamically adding and removing paths or adapters

When ESS devices are configured as MPIO-capable devices under AIX 5.2.0 ML3, you can add or remove extra paths or adapters while I/O is running. To add extra paths that are attached to an adapter to existing available devices, enter:

mkpath -l hdiskX -p fscsiY When the command returns successfully, the paths are added to the devices.

To check the device configuration status, enter:

lspath -l hdiskX

or

pcmpath query device X

To add a new adapter to existing available ESS MPIO devices, enter:

cfgmgr -vl fscsiX

To check the adapter configuration status, enter:

pcmpath query adapter

or

pcmpath query device

To dynamically remove all paths under a parent adapter from an MPIO device, enter:

rmpath -dl hdiskX -p fscsiY

To dynamically remove an adapter and all children devices from ESS MPIO devices, use **smit mpio**, or enter the following on the command line:

rmdev -l fscsiX -R

or

rmdev -dl fscsiX -R

Note: You cannot remove last path from an ESS MPIO device. The command will fail if you try to remove the last path from an ESS MPIO device.

Removing SDDPCM from an AIX host system

Before you remove the SDDPCM package from your AIX host system, all ESS 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: SDDPCM supports MPIO-capable ESS devices as the boot device. If your system has ESS boot devices configured with SDDPCM, then there is no method available for you to migrate ESS boot device from SDDPCM to the AIX default PCM.

After all the ESS devices are removed or migrated to the AIX default PCM and the SDDPCM server daemon (pcmsrv) is stopped, perform 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. Enter the following installation package names in the **SOFTWARE name** field: **devices.sddpcm.52b.rte devices.fcp.disk.ibm2105.mpio.rte** and then press **Enter**.

   Alternatively, you can 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.

   Note: To remove SDDPCM, you must remove both the ESS host attachment for SDDPCM and the SDDPCM software packages before you reconfigure ESS devices or reboot the system. Otherwise, the devices can be in the *Defined* state and will not be able to be configured as either MPIO or non-MPIO devices.
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 ESS 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 ESS host attachment for SDDPCM are removed from your system.

Using SDDPCM pcmpath commands

SDDPCM supports the following **pcmpath** commands:

- **pcmpath query adapter [n]**
- **pcmpath query adaptstats [n]**
- **pcmpath query device [n]**
- **pcmpath query devstats [n]**
- **pcmpath set adapter n online | offline**
- **pcmpath set device M path N online | offline**
- **pcmpath disable port <location> ess <essid>**
- **pcmpath enable port <location> ess <essid>**
- **pcmpath open device <m> path <n>**
- **pcmpath query essmap**
- **pcmpath query portmap**
- **pcmpath query wwpn**

Note: If the commands are used for a device, then 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, then 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.

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 ESS MPIO devices or paths to an ESS 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 an ESS port or ports to online or offline. This section includes descriptions of these commands. Table 20 provides an alphabetical list of these commands, a brief description, and where to go in this chapter for more information.

Table 20. Commands

Command	Description	Page
pcmpath disable ports	Places paths connected to certain ports offline.	119
pcmpath enable ports	Places paths connected to certain ports online.	119
pcmpath open device path	Opens an INVALID path.	123
pcmpath query adapter	Displays information about adapters.	125
pcmpath query adaptstats	Displays performance information for all FCS adapters that are attached to SDDPCM devices.	126
pcmpath query device	Displays information about devices.	127
pcmpath query devstats	Displays performance information for a single SDDPCM device or all SDDPCM devices.	129
pcmpath query essmap	Displays each device, path, location, and attributes.	131

Table 20. Commands (continued)

Command	Description	Page
pcmpath query portmap	Displays ESS MPIO device port location.	132
pcmpath query wwpn	Displays the World Wide Port Name (WWPN) for all fibre-channel adapters.	133
pcmpath set adapter	Sets all device paths that are attached to an adapter to online or offline.	134
pcmpath set device path	Sets the path of a device to online or offline.	135

pcmpath disable ports

The **pcmpath disable ports** command sets MPIO device paths offline for specified ESS location code.

Syntax

Parameters

location

The location 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

essid

The ESS 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                 fscsi1/path2            CLOSE  OFFLINE          11          0
    3                 fscsi1/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                 fscsi1/path2            CLOSE  OFFLINE       8816          0
    3                 fscsi1/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               fscsi1/path2            CLOSE  OFFLINE       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  OFFLINE       9044          0
    1                 fscsi0/path1            CLOSE   NORMAL       9084          0
    2                 fscsi1/path2            CLOSE  OFFLINE       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  OFFLINE       9089          0
    1                fscsi0/path1            CLOSE   NORMAL       9238          0
    2                fscsi1/path2            CLOSE  OFFLINE       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  OFFLINE       9059          0
    1               fscsi0/path1              CLOSE   NORMAL       9121          0
    2               fscsi1/path2              CLOSE  OFFLINE       9143          0
    3               fscsi1/path3              CLOSE   NORMAL       9073          0
```

pcmpath enable ports

The **pcmpath enable ports** command sets MPIO device paths online for the specified ESS location code.

Syntax

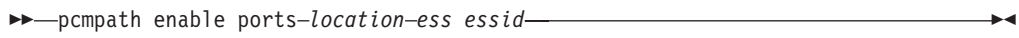

Parameters

location

The location 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

essid

The ESS 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                 fscsi1/path2           CLOSE   NORMAL         11           0
    3                 fscsi1/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                 fscsi1/path2           CLOSE   NORMAL       8816          0
    3                 fscsi1/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 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 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 is INVALID.

If 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: load balancing, round robin, or failover.

Serial TheLUN for this device.

Path# The path id displayed by the **pcmpath 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 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 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 input and output.

Errors The number of input and output 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

```
►►—pcmpath query adapter—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.

Examples

If you enter the **pcmpath query adapter** command, the following output is displayed:

```
Active Adapters :2

Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active
    0           fscsi2  NORMAL   ACTIVE  920506           0      80       38
    1           fscsi0  NORMAL   ACTIVE  921100           0      80       38
```

The terms used in the output are defined as follows:

Adpt #
The index number of the adapter.

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 input or output.

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 a device number, information about all devices 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.

Examples

If you enter the **pcmpath query adaptstats 0** command, the following output is displayed:

```
Adapter #:  0
=============
                 Total Read  Total Write  Active Read  Active Write   Maximum
I/O:               1105909           78            3             0        11
SECTOR:            8845752            0           24             0        88
Adapter #:  1
=============
                Total Read  Total Write  Active Read  Active Write   Maximum
I/O:                 1442           78            3             0        11
SECTOR:            156209            0           24             0        88


/*--------------------------------------------------------------------------*/
```

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 MPIO device or all MPIO devices. If you do not enter a device number, information about all devices is displayed. If a device number is entered, then the command will display the device information about the hdisk that is associated with this number.

The **pcmpath query device** commands displays only ESS MPIO-capable devices that are configured with the SDDPCM module. Any AIX internal disks or non-SDDPCM-configured ESS MPIO-capable devices will not be displayed.

Syntax

```
►►—pcmpath query device— device number————————————————————————————►◄
```

Parameters

device number
The device number refers to the logical device number of the hdisk.

Examples

If you enter the **pcmpath query device 2** command, the following output about hdisk2 is displayed:

For ESS:

```
DEV#:   2  DEVICE NAME: hdisk2  TYPE: 2105800  ALGORITHM:  Load Balance
SERIAL: 00923922
==========================================================================
Path#       Adapter/Path Name         State     Mode     Select     Errors
    0           fscsi0/path0          CLOSE   NORMAL          0          0
    1           fscsi0/path1          CLOSE   NORMAL          0          0
    2           fscsi1/path2          CLOSE   NORMAL          0          0
    3           fscsi1/path3          CLOSE   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, round robin, or failover

Serial The LUN for this device.

Path The path ID.

Adapter
The name of the adapter to which the path is attached.

Path Name
The name of the path.

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 was removed by SDDPCM 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 failed to open, while the device did 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 that occurred on a path of this device.

pcmpath query devstats

The **pcmpath query devstats** command displays performance information for a single MPIO device or all MPIO devices. If you do not enter a device number, information about all devices is displayed. If a device number is entered, then the command will display the device information about the hdisk that is associated with this number.

The **pcmpath query devstats** command displays only MPIO-capable devices that have been configured with the SDDPCM module. Any AIX internal disks or non-SDDPCM-configured MPIO-capable devices will not be displayed.

Syntax

```
►►—pcmpath query devstats—— device number————————————————————————►◄
```

Parameters

device number
The device number refers to the logical device number of the hdisk.

Examples

If you enter the **pcmpath query devstats 2** command, the following output about hdisk2 is displayed:

```
DEV#:   2  DEVICE NAME: hdisk2
===============================
                Total Read  Total Write  Active Read  Active Write   Maximum
I/O:                 60           10            0             0         2
SECTOR:             320            0            0             0        16

Transfer Size:      <= 512        <= 4k       <= 16K        <= 64K     > 64K
                        30           40            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

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 fcppcmmap program invoked at system reboot phase, or by the pcmpath program if any device configuration change occurs after system rebooted. fcppcmmap issues SCSI commands directly to the disks (inquiry, read capacity, and log sense) to collect the information displayed.

Syntax

```
►►—pcmpath query essmap————————————————————————————————————————————►◄
```

Examples

If you enter the **pcmpath query essmap** command, the following output is displayed:

```
 Disk      Path        Location     adapter   LUN SN       Type             Size   LSS    Vol   Rank  C/A   S   Connection    port  RaidMode
-------   ------      ---------    --------  -------      ----             ----   ---    ---   ----  ---   -   ----------    ----  --------
hdisk3     path0      10-68-01[FC] fscsi0   20112028  IBM 2105-E20          1.0    18     1    1201   02   Y   R1-B1-H3-ZA     8          5
hdisk3     path1      10-68-01[FC] fscsi0   20112028  IBM 2105-E20          1.0    18     1    1201   02   Y   R1-B3-H4-ZA    8c          5
hdisk3     path2      20-58-01[FC] fscsi1   20112028  IBM 2105-E20          1.0    18     1    1201   02   Y   R1-B1-H3-ZA     8          5
hdisk3     path3      20-58-01[FC] fscsi1   20112028  IBM 2105-E20          1.0    18     1    1201   02   Y   R1-B3-H4-ZA    8c          5
hdisk4     path0      10-68-01[FC] fscsi0   20712028  IBM 2105-E20          1.0    18     7    1201   02   Y   R1-B1-H3-ZA     8          5
hdisk4     path1      10-68-01[FC] fscsi0   20712028  IBM 2105-E20          1.0    18     7    1201   02   Y   R1-B3-H4-ZA    8c          5
hdisk4     path2      20-58-01[FC] fscsi1   20712028  IBM 2105-E20          1.0    18     7    1201   02   Y   R1-B1-H3-ZA     8          5
hdisk4     path3      20-58-01[FC] fscsi1   20712028  IBM 2105-E20          1.0    18     7    1201   02   Y   R1-B3-H4-ZA    8c          5
```

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.

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 ESS.

Type The device and model.

Size The configured capacity of the LUN in GB.

LSS The logical subsystem where the LUN resides.

Vol The volume number within the LSS.

Rank The unique identifier for each RAID array within the ESS.

C/A The cluster and adapter accessing the array.

S Indicates that the device is shared by two and more ESS ports. Valid values are *yes* or *no*.

Connection The physical location code of ESS adapter through which the LUN is accessed.

Port The ESS port through which the LUN is accessed.

RaidMode The disk RAID mode.

pcmpath query portmap

The **pcmpath query portmap** command displays hdisk status from an ESS physical location code perspective.

Syntax

Examples

If you enter the **pcmpath 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
                  A B A B A B A B    A B A B A B A B    A B A B A B A B    A B A B A B A B
12028    hdisk3    - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
12028    hdisk4    - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
12028    hdisk5    - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
12028    hdisk6    - - - - - - - -    N - - - - - - -    - - - - - - - -    0 - - - - - - PD
12028    hdisk7    - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
12028    hdisk8    - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
12028    hdisk9    - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
```

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.

O 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.

- The path is not configured.

PD The path is down. It is either not functional or has been placed offline.

pcmpath query wwpn

The **pcmpath query wwpn** command displays the World Wide Port Name (WWPN) of the host fibre-channel adapters.

Syntax

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.

Syntax

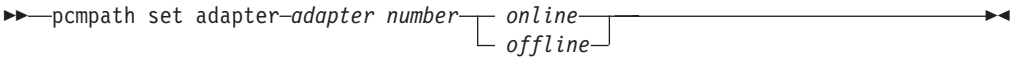

Parameters

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.

Examples

If you enter the **pcmpath set adapter 0 offline** command:

- Adapter 0 changes to *Offline* mode and, if there are some paths in the opened state, its state might change to *failed*.
- All paths attached to adapter 0 change to *Offline* mode and their states change to Dead, if they were in the Open state.

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.

Syntax

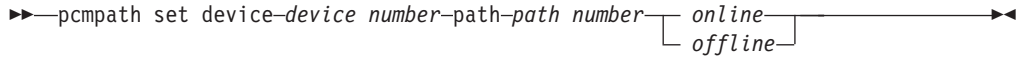

Parameters

device number
The logical device number of the hdisk.

path number
The path id that you want to change.

online
Enables the path for service.

offline
Disables the path from service.

Examples

If you enter the **pcmpath set device 0 path 0 offline** command, path 0 for device 0 changes to Offline mode.

SDDPCM server daemon

The SDDPCM server (also referred to as *pcmsrv*) is an integrated component of SDDPCM 2.0.1.0 (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. For detailed information about the SDDPCM server daemon, see Chapter 11, “Using the SDD server and the SDDPCM server,” on page 279.

Verifying if the SDDPCM server has started

After you have installed SDDPCM and rebooted 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 ESS MPIO devices have not been configured yet. In this case, you can either reboot the system or you can start pcmsrv manually after ESS MPIO devices are configured.

Because pcmsrv is bound to SDDPCM, 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 ESS MPIO devices are configured.

Starting the SDDPCM server manually

If pcmsrv did not start automatically after you performed the SDDPCM Installation and configured ESS devices, you can start pcmsrv by entering **startsrc -s pcmsrv**.

Go to “Verifying if the SDDPCM server has started” to see if you successfully started the SDDPCM server.

Stopping the SDDPCM server

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

SDDPCM server (pcmsrv) recovers paths for SDDPCM which are not recovered by the SDDPCM healthchecker. You should not stop pcmsrv permanently. However, in certain cases when you want to permanently (not to start even after a system reboot) disable pcmsrv, you need to comment out the following line in the system init table (/etc/inittab):

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

See Chapter 11, “Using the SDD server and the SDDPCM server,” on page 279 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 281)
- Changing to a different port number for the SDDPCM server (“Changing the sddsrv or pcmsrv TCP/IP port number” on page 282)

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.

AIX 5.2.0 ML3 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

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 to fail a broken path. By default, this attribute has a value of *delayed_fail*, which means that the fast failure is disabled.

To enable fast failure, you must do the following:

1. Change all the children devices of fscsiX on your system to the defined state by running the command **’rmdev -l fscsiX -R’** for all adapters on your system
2. Enter the command **’chdev -l fscsiX -a fc_err_recov=fast_fail’** for all adapters on your system.
3. Enter **cfgmgr** 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

Fibre channel dynamic device tracking

Dynamic Tracking support allows for the following scenarios without I/O errors or loss of access:

1. Combine two switches in two different SANS into one SAN by connecting the switches with a cable and cascading the switches.
2. Change the ESS FC cable from one switch port to a different switch port; Disconnected cable must be reconnected within 15 seconds.
3. Swap two ESS FC cables on the switch ports. The disconnected cables must be reconnected within 15 seconds.
4. Swap two ESS FC cables on the ESS ports. The disconnected cables must be reconnected within 15 seconds.

Note: Fibre channel device dynamic tracking does not support the case when a cable is moved from one adapter on the ESS to another free, previously-unseen adapter on the ESS. The reason is that the World Wide Port Name (WWPN) will be different for that adapter, and the ESS 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 children devices of fscsiX on your system to the defined state by executing the **'rmdev -l fscsiX -R'** command for all adapters on your system.
2. Execute the**'chdev -l fscsiX -a dyntrk=yes'** command for all adapters on your system.
3. Execute **cfgmgr** to reconfigure all devices back to the available state.

For more information about Dynamic Tracking and Fast I/O Failure, refer to the AIX release notes at:

http://publib.boulder.ibm.com/pseries/aixgen/relnotes/52RELNOTES/10073902.htm

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.

Chapter 4. Using SDD on a 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. SDD 1.4.0.4 supports the coexistence of ESS devices and SAN Volume Controller devices. SDD 1.5.0.4 (or later) supports the coexistence of ESS, SAN Volume Controller, and SAN Volume Controller for Cisco MDS 9000 devices.

For updated and additional information that is not included in this manual, please see the Readme file on the CD-ROM or go to SDD Web site at:

www-1.ibm.com/servers/storage/support/software/sdd.html

Verifying the hardware and software requirements

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

Hardware

- 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 vpaths.

  For information on the fibre-channel adapters that can be used on your HP-UX host system go to:

  www-1.ibm.com/servers/storage/support/software/sdd.html
- A SCSI cable to connect each SCSI host adapter to a storage system controller port
- Subsystem LUNs that have been created and confirmed for multiport access
- A fiber-optic cable to connect each fibre-channel adapter to a supported storage device port

Software

A PA-RISC system running HP-UX 11.0 or HP-UX 11i (fibre-channel only)

Unsupported environments

SDD does not support the following environments:

- 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 and installation of LIC nor during any ESS concurrent maintenance that impacts the path attachment, such as an ESS-host-bay-adapter replacement
- Single-path configuration for Fibre Channel

© Copyright IBM Corp. 1999, 2004

SDD support for 32-bit and 64-bit applications on HP-UX 11.0 and HP-UX 11i (fibre-channel)

SDD supports 32-bit and 64-bit applications on HP-UX 11.0 and HP-UX 11i (Fibre Channel).

Attention: For SDD to operate properly, ensure that the patches listed in Table 22 on page 143 or newer patches are installed on your HP-UX 11.0 host system.

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 1.

Preparing for SDD installation

Before you install SDD, you must configure the supported storage devices to your host system and attach the required SCSI or fibre-channel adapters.

Configuring the ESS

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, refer to the *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide*.

Configuring the SAN Volume Controller

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.

For information about configuring your SAN Volume Controller, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Configuration Guide*.

Configuring the SAN Volume Controller for Cisco MDS 9000

Before you install SDD, configure your SAN Volume Controller for Cisco MDS 9000 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 for Cisco MDS 9000, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Configuration Guide*.

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 21 further describes the various installation scenarios and how you should proceed.

Table 21. SDD installation scenarios

Installation scenario	Description	How to proceed
Scenario 1	• SDD is not installed. • The SDD server for Expert is installed. • No software application or DBMS communicates directly to sdisk interface.	Go to: 1. "Determining if the SDD 1.3.1.5 (or later) server for Expert is installed" on page 144 2. "Installing SDD" on page 144 3. "Standard UNIX applications" on page 160
Scenario 2	• 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: 1. "Determining if the SDD 1.3.1.5 (or later) server for Expert is installed" on page 144 2. "Installing SDD" on page 144 3. "Using applications with SDD" on page 160
Scenario 3	• SDD is installed. • The SDD server for Expert is installed.	Go to 1. "Determining if the SDD 1.3.1.5 (or later) server for Expert is installed" on page 144 2. "Upgrading SDD" on page 149

For SDD to operate properly, ensure that the patches shown in Table 22 are installed on your HP-UX 11.0 host system.

Note: Do *not* install the patches noted in Table 22 on HP-UX 11i (fibre channel) host systems.

Table 22. Patches necessary for proper operation of SDD on HP-UX 11.0

Application mode	HP-UX patch	Patch description
32-bit	PHKL_20674	Fix VxFS unmount hang and NMF, sync panics
32-bit	PHKL_20915	Trap-related panics or hangs
32-bit	PHKL_21834	Fibre-channel Mass Storage Driver patch
32-bit	PHKL_22759	SCSI I/O Subsystem Cumulative patch
32-bit	PHKL_23001	Signal, threads, spinlock, scheduler, IDS, q3p
32-bit	PHKL_23406	Probe, sysproc, shmem, thread cumulative patch
32-bit or 64-bit	PHKL_21392	VxFS performance, hang, icache, DPFs
32-bit or 64-bit	PHKL_21624	Start, JFS, PA8600, 3Gdata, NFS, IDS, PM, VM, async
32-bit or 64-bit	PHKL_21989	SCSI IO Subsystem Cumulative patch

Table 22. Patches necessary for proper operation of SDD on HP-UX 11.0 (continued)

Application mode	HP-UX patch	Patch description
64-bit	PHKL_21381	Fibre-channel Mass Storage Driver
32-bit	PHCO_29436	som2elf cumulative patch

Determining if the SDD 1.3.1.5 (or later) server for Expert is installed

If you previously installed the SDD server (the standalone version) for IBM TotalStorage Expert V2R1 (ESS Expert) on your HP-UX host system, you must remove this standalone 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 standalone version of the SDD server (for ESS Expert).

To determine if the standalone version of the SDD server is installed on your host system, enter:

swlist SDDsrv

If you previously installed the standalone version of the SDD server, the output from the **swlist SDDsrv** command looks similar to this:

```
SDDsrv  1.0.0.0  SDDsrv bb-bit Version: 1.0.0.0 Nov-14-2001 15:34
```

Note:

- The installation package for the standalone 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 standalone SDD Server installation package is SDDsrvHP32_020115.depot for a 32-bit environment, and SDDsrvHP64_020115.depot for a 64-bit environment.)
- For instructions on how to remove the standalone 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 Web site:

  www-1.ibm.com/servers/storage/support/software/swexpert.html

For more information about the SDD server daemon, go to “SDD server daemon” on page 155.

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.

Perform 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 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 IBMsdd.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*/hp32bit/IBMsdd.depot
   c. Click **OK**.
      You will see output similar to that in either Figure 4 or Figure 5.

```
Name           Revision        Information                    Size(Kb)
IBMsdd_tag ->  B.11.00.01      IBMsdd Driver 32-bit           nnnn
```

Figure 4. IBMsdd Driver 32-bit

```
Name           Revision        Information                    Size(Kb)
IBMsdd_tag ->  B.11.00.01      IBMsdd Driver 64-bit           nnnn
```

Figure 5. IBMsdd Driver 64-bit

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          : Building kernel
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 will 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 executes successfully.

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 the following tables.

Table 23. SDD components installed for HP-UX host systems

File	Location	Description
mod.o	/opt/IBMsdd/bin	Object file for SDD driver module
Executables	/opt/IBMsdd/bin	Configuration and status tools
README.sd	/opt/IBMsdd	README file
defvpath	/sbin	SDD configuration file used during startup
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 defvpath 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 vpath filesystems

Table 23. SDD components installed for HP-UX host systems (continued)

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
querysn.1	/usr/local/man/man1/querysn.1	Manpage for querysn
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

Table 24. 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 25. SDD commands and their descriptions for HP-UX host systems

Command	Description
cfgvpath [-c]	Configures vpath devices. Updates the information in /etc/vpath.cfg and /etc/vpathsave.cfg. It does not change the running system. It sets up the system for reboot. **Note: cfgvpath -c** updates the configuration file but does not update the kernel. To update the kernel, you need to reboot.
cfgvpath [-r] (Dynamic reconfiguration)	The vpath configuration is updated without system reboot. If initially there is no vpath configured, **cfgvpath -r** will fail with message "failed to get information from kernel, don't run dynamic configuration, do cfgvpath instead." In this case, execute **cfgvpath** without any option.
defvpath	Second part of the **cfgvpath** command configuration during startup time.
showvpath	Lists the configuration mapping between SDD devices and underlying disks.
datapath	SDD driver console command tool.
hd2vp	Converts a volume group from sdisks into SDD vpaths.
vp2hd	Converts a volume group from SDD vpaths into sdisks.
rmvpath [-all, -vpathname]	Removes vpath devices from the configuration.
gettrace	Debug tool that gets trace information when problem happens.

Table 25. SDD commands and their descriptions for HP-UX host systems (continued)

Command	Description
man	Manpage for SDD commands, for example, man datapath. Supported SDD commands are datapath, gettrace, hd2vp, querysn, rmvpath, sddsrv, showvpath, and vp2hd.

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 SDD. Go to "Standard UNIX applications" on page 160 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 160 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/defvpath
- /opt/IBMsdd/bin/showvpath
- /opt/IBMsdd/bin/master
- /opt/IBMsdd/bin/system
- /opt/IBMsdd/bin/mod.o
- /opt/IBMsdd/bin/rmvpath
- /opt/IBMsdd/bin/ckvpath
- /opt/IBMsdd/bin/get_root_disks
- /opt/IBMsdd/bin/gettrace
- /opt/IBMsdd/bin/hd2vp
- /opt/IBMsdd/bin/vp2hd
- /sbin/defvpath
- /sbin/cfgvpath
- /sbin/datapath
- /sbin/get_root_disks
- /sbin/rmvpath
- /sbin/showvpath
- /sbin/hd2vp
- /sbin/vp2hd
- /sbin/sddsrv

- /etc/sample_sddsrv.conf

During installation, /opt/IBMsdd/bin/cfgvpath program is initiated to create vpaths in the /dev/dsk and /dev/rdsk directories for all IBM disks that are available on the system. After installation is done, all vpaths are configured and the driver is loaded. The system will not reboot.

Note: SDD devices are found in /dev/rdsk and /dev/dsk. The device is named according to the SDD number. A device with a number of 0 would be /dev/rdsk/vpath1.

Upgrading SDD

Upgrading SDD consists of removing and reinstalling the IBMsdd package. If you are upgrading SDD, go to “Uninstalling the Subsystem Device Driver” on page 150 and then go to “Installing SDD” on page 144.

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. Execute **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 vpaths to volume groups containing cXtXdX devices.
5. When the volume groups have been converted:
   a. Remove the /etc/vpathsave.cfg file.
   b. Execute **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 vpath devices looks good in the **datapath query device** output, use the **hd2vp** script to change volume groups to volume groups containing SDD vpaths.

   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.

Uninstalling the Subsystem Device Driver

The following procedure explains how to remove the 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 vpaths 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 144.

Configuring SDD

This section provides information necessary to configure SDD. Use the HP Command Line Interface (CLI) to manage SDD devices.

Changing an SDD hardware configuration

When adding or removing multiport SCSI devices, you must reconfigure SDD to recognize the new devices. Perform the following steps to reconfigure SDD:

1. Issue the **cfgvpath** command to reconfigure the vpath by entering:

   /opt/IBMsdd/bin/cfgvpath -c

2. Restart the system by entering:

   shutdown -r 0

The **querysn** command can be used to list all ESS devices visible to the host. The **querysn** command reads the unique serial number of an ESS 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 two conversion scripts:

hd2vp The hd2vp script converts a volume group from supported storage device sdisks into SDD vpaths.

The syntax for hd2vp script is as follows:

```
hd2vp vgname
```

vp2hd The vp2hd script converts a volume group from SDD vpaths 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 vgname
```

These two conversion programs require that a volume group contain either *all* original supported storage device sdisks or *all* SDD vpaths. The program fails if a volume group contains both kinds of device special files (mixed volume group).

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 vpath configuration in memory and then identifies a list of differences. It then issues commands to update the vpath configuration in memory with the current hardware configuration. The commands that **cfgvpath -r** issues to the vpath driver are:

   a. Add vpath device.
   b. Remove a vpath device; this will fail if device is busy.
   c. Add path to vpath device.
   d. Remove path for a vpath device; this will fail deletion of path if device is busy, but will set path to DEAD and OFFLINE.

2. **rmvpath** command removes one or more vpath devices.

```
rmvpath -all                 # Remove all vpaths

rmvpath vpath_name           # Remove one vpath 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 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 303 for more information about the **datapath set device policy** command.

Preferred node path-selection algorithm for the SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000

The SAN Volume Controller and the SAN Volume Controller for Cisco MDS 9000 is a two-controller disk subsystem. SDD distinguishes the paths to a SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 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, then 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 12, "Using the datapath commands," on page 283 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 execute 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 and SAN Volume Controller for Cisco MDS 9000 Serials are 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
```

Example output from **datapath query device** command issued in SDD 1.4.0.0 (or later): (This example shows a SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 device and an ESS device.)

```
Dev#:   2     Device Name: vpath4  Type: 2145       Policy:    Optimized
 Serial: 60056768018506870000000000000000
============================================================================
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.

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 11, “Using the SDD server and the SDDPCM server,” on page 279 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 sddsrv has *not* started, you should uninstall SDD and then reinstall SDD. See “Installing SDD” on page 144 for more information.

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 `nohup /sbin/sddsrv&`

Go to “Verifying if the SDD server has started” for the steps to see if you successfully started the SDD server.

Changing to a different port number for the SDD server

See “Changing the sddsrv or pcmsrv TCP/IP port number” on page 282.

Stopping the SDD server

Perform the following steps to stop the SDD server:

1. Query the process for SDD server by entering the following command:

   ps -ef | grep sddsrv
2. Find and record the process ID for /sbin/sddsrv
3. Stop the process ID (that you recorded in step 2) by entering the following command:

   kill -9 *pid*

   where *pid* represents the process ID that you recorded in step 2.

Understanding the SDD 1.3.1.5 (or later) support for single-path configuration for ESS

SDD 1.3.1.5 (or later) does not support concurrent download and installation of the ESS licensed internal code (LIC) on single-path mode.

However, SDD supports single-path SCSI or fibre-channel connection from your HP-UX host system to an ESS. It is possible to create a volume group or a vpath device with only a single path.

Notes:

1. SDD supports one SCSI adapter on the host system.
2. If your host has only one fibre-channel adapter, it requires you to connect through a switch to multiple ESS ports. You should have at least two fibre-channel adapters to prevent data loss due to adapter hardware failure or software failure.
3. Because of single-path connection, SDD can not provide single-point-failure protection and load balancing. IBM does not recommend using a single-path connection.

The policy will be default to SINGLE_PATH_POLICY.

You can change from single-path to multipath configuration by using the **cfgvpath -r** command. For more information about the **cfgvpath -r** command, see “Dynamic reconfiguration” on page 151.

Understanding the SDD error recovery policy

The new SDD error recovery policy (1.3.2.2 or higher) differs from the error recovery policy of SDD levels 1.3.2.1 or earlier. The new error recovery policy reports errors to applications more quickly, preventing unnecessary retries.

SDD 1.3.2.1 (or earlier) error recovery policy

With SDD 1.3.2.1 (or earlier), error recovery policy is designed to cover a transient kind of error from the user applications. The error recovery policy prevents a path from becoming disabled in the event of transient errors. This policy can cause longer interruption for recovery to take place. The recovery policy halts the I/O activities on functional paths and vpath devices for some time before failing paths are set to DEAD state.

SDD 1.3.2.2 (or later) error recovery policy

With SDD 1.3.2.2 (or later), error recovery policy is designed to report failed I/O requests to applications more quickly. This process prevents unnecessary retries, which can cause the I/O activities on good paths and vpath devices to halt for an unacceptable period of time.

SDD assumes the following two modes of operation:

- Single-path mode (for ESS only)
- Multiple-path mode

Single-path mode (for ESS only)

In single-path mode, only a single path is available in access to an ESS LUN.

SDD error recovery policy in single-path mode has the following characteristics:

- When an I/O error occurs, SDD retries the I/O operation up to two times.
- With SDD 1.3.2.2 (or later), SDD returns a failed I/O to the application and sets the state of this failing path to DEAD. The SDD driver relies on the SDD server daemon to detect the recovery of the single path. The SDD server daemon recovers this failing path and changes its state to OPEN. (SDD can change a single and failing path into DEAD state.)
- With the SDD 1.3.2.1 (or earlier) error recovery policy, SDD returns the failed I/O to the application and leaves this path in OPEN state (SDD never puts this single path into DEAD state).
- With SDD 1.3.2.2 (or later), when the SDD server daemon detects that the single CLOSE path is failing, it will change the state of this failing path to

CLOSE_DEAD. When the SDD server daemon detects that a CLOSE_DEAD path has recovered, it will change the state of this path to CLOSE. With a single path configured, the SDD vpath device can not be opened if the only path is in a CLOSE_DEAD state.

Multipath mode

Both the SDD 1.3.2.2 (or later) and SDD 1.3.2.1 (or earlier) error recovery policies in multiple-path mode have the following common characteristics:

- If an I/O error occurs on the last operational path to a device, SDD attempts to reuse (or fail back to) a previously failed path.

The host system has multiple paths that are configured to an ESS LUN. SDD 1.3.2.2 (or later) error recovery policy in multiple-path mode has the following characteristics:

- If an I/O error occurs on a path, SDD attempts to use the path again after three successful I/O operations on another operational path.
- If the consecutive error count on the path reaches three, SDD puts the path into the DEAD state.
- If an I/O error occurs on the last operational path to a device, SDD attempts to reuse (or perform a failback operation to return to) a previously-failed path.
- If the SDD server daemon detects that the last CLOSE path is failing, the daemon will change the state of this path to CLOSE_DEAD. The SDD server can automatically recover the path if it is detected that it is functional.
- If an I/O fails on all open paths to an ESS LUN, SDD returns the failed I/O to the application and changes the state of all open paths (for failed I/Os) to DEAD, even if the consecutive error count on some of the paths did not reach three.
- If an open path already failed some I/Os, it will not be selected as a retry path.

The SDD 1.3.2.1 (or earlier) error recovery policy in multiple-path mode has the following characteristics:

- If an I/O error occurs on a path, SDD 1.3.2.1 does not attempt to use the path until 2000 successful I/O operations occur on an operational path.
- The last path is reserved in OPEN state.
- If an I/O fails on all OPEN paths to an ESS LUN, SDD returns the failed I/O to the application and leaves all the paths in OPEN state.
- A failed I/O is retried on all OPEN paths to an ESS LUN even if the OPEN path already failed I/Os.
- SDD changes the failed path from the DEAD state back to the OPEN state after 50 000 successful I/O operations on an operational path.

Note: You can display the error count with the **datapath** command.

Dynamically opening an Invalid or CLOSE_DEAD path

With SDD 1.3.2.0 (or later), you can issue the **datapath open path** command to dynamically open a path that is in an INVALID or CLOSE_DEAD state. You can use this command even when the I/O is actively running. See “datapath open device path” on page 286 for more information.

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 vpaths 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 perform the following tasks:

- 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."
- FTP the map file to the other node within a high availability cluster. See "Moving the map file."
- Create the volume group device directory. See "Creating the volume group device directory" on page 159.
- Create the group special file. See "Creating the group special file" on page 159.

For more information about the **vgimport** command, see "Importing volume groups" on page 159.

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 the –s option is specified, then the –p, –v, and –m options must also be specified. 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 vgvpath1*

where */tmp/vgpath1.map* represents your mapfile, and *vgvpath1* 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 *vgvpath1.map*map file to node 2, enter:

rcp */tmp/vgvpath1.map node2:/tmp/vgvpath1.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/vgvpath1* on node 2, enter:

mkdir */dev/vgvpath1*

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/vgvpath1/group c 64 *n*

where *n* is the same as that was given when */dev/vgvpath1/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 (ESS Series 800 only). When the –s option is specified, then the –p, –v, and –m options must also be specified. 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 vgvpath1*

where */tmp/vgpath1.map* represents your mapfile, and *vgvpath1* represents the path name of the volume group that you want to import.

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 DBMS for use with SDD:

- Standard UNIX applications
- Network File System (NFS) file server
- Oracle

Standard UNIX applications

If you have not already done so, install SDD using the procedure in “Installing SDD” on page 144. 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 devices and their underlying disks.

To use the SDD driver for an existing logical volume, you must remove the existing logical volume and volume group and re-create it using the SDD device.

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 super-user privileges to perform 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          lv              lvm
```

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.

   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/rdsk/vpath1

      A message similar to the following is displayed:

      ```
      Physical volume "/dev/rdsk/vpath1" has been successfully created.
      ```

      In this example, the SDD 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/rdsk/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**, see the following Web site:

   http://docs.hp.com/

   Click *HP-UX Reference (Manpages)*. Then refer to *HP-UX Reference Volume 2*.

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 lvol1:

   # umount /dev/vgibm/lvol1

   Next, remove the logical volume.

   For example, enter the following command to remove logical volume lvol1:

   # lvremove /dev/vgibm/lvol1

   A message similar to the following is displayed:

```
The logical volume "/dev/vgibm/lvol1" 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/lvol1" 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:

   # vgremove /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 super-user privileges to perform these subtasks.

As an example, suppose you have a logical volume called lvol1 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 lvol1 to use SDD.

1. Determine the size of the logical volume.

   Enter the following command to determine the size of the logical volume:

   # lvdisplay /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/rdsk/vpath1

   A message similar to the following is displayed:

```
Physical volume "/dev/rdsk/vpath1" has been successfully created.
```

   In this example, the SDD 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/rlvol1".
      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 correctly set for SDD to operate properly. This is particularly true if you are going to be using concurrent microcode download.

   If you are going to be using concurrent microcode download with multipath SCSI, perform the following steps to set the proper timeout value for the logical volume manager:

   a. Ensure that the timeout value for an SDD logical volume is set to the default. Enter **lvdisplay /dev/vgibm/lvoly** and press **Enter**. If the timeout value is not default, enter **lvchange -t 0 /dev/vgibm/lvoly** and press **Enter** to change it. (In this example, vgibm is the name of the logical volume group that was previously configured to use SDD; in your environment the name may be different.)
   b. Change the timeout value for an SDD physical volume to 240. Enter **pvchange -t 240 /dev/dsk/vpathn** and press **Enter**. (*n* refers to the vpath number.) If you are not sure about the vpath number, enter **/opt/IBMsdd/bin/showvpath** and press **Enter** to obtain this information.
   c. The re-created logical volume must be mounted before it can be accessed.

Note: During a concurrent code download (CCL) of licensed internal code, certain types of fabric errors, or a failure of an SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 cluster node, the remaining node in the IOGroup temporarily take remedial actions to protect customer data. The latency of the host I/O to the SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 can increase to 60 seconds or more. Because the default setting of the HP disk device driver timeout value is 30 seconds, path failures on the remaining node can result and SDD will have no more paths available The timeout value of the physical device should be changed for all SDD vpath devices using the **pvchange** command. You should do this after the physical volume has been created (using **pvcreate**) and added to a volume group (using **vgcreate**). For example, enter **pvchange -t 90 /dev/dsk/vpath[#]**.

For additional information about SAN Volume Controller, refer to the following Web site, click **Technical Notes**, and browse for more information:

www-1.ibm.com/servers/storage/support/virtual/2145.html

For additional information about SAN Volume Controller for Cisco MDS 9000, refer to the following Web site, click **Technical Notes**, and browse for more information:

www-1.ibm.com/servers/storage/support/virtual/2062-2300.html

In some cases it might be necessary to use standard HP 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**, see the following Web site:

http://docs.hp.com/

Click *HP-UX Reference (Manpages)*. Then, refer to *HP-UX Reference Volume 2*.

Installing SDD on a Network File System file server

The procedures in this section show how to install SDD for use with an exported file system (Network File System file server).

Setting up Network File System for the first time

Perform 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 144.
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/rdsk/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. Click **yes** in the **mount at boot** field.

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 Network File System file server

Perform the following steps if you have the Network File System 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:
   - Passes the start time **fsck pass**
   - Mounts properly
   - 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 Network File System service. Then review your steps and try again.

Oracle

You must have super-user privileges to perform the following procedures. You also need to have Oracle documentation available to use. 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 differs depending on the choice that you make.

Using a file system:

1. If you have not already done so, install SDD using the procedure in “Installing SDD” on page 144.
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 installing 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: When using raw partitions, 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. Make sure that all the databases are closed before making changes.

In the following procedure you will be replacing the raw devices with the SDD devices.

1. If you have not already done so, install SDD using the procedure in "Installing SDD" on page 144.
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 the installation of Oracle8 on a file system residing on an SDD partition.
   b. Set up the Oracle user's ORACLE_BASE and ORACLE_ HOME environment variables to the 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 *Oracle8 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 (files named /dev/rdsk/vpathNs, where *N* is the SDD instance number, and *s* is the partition ID) that point to the slice.
3. Determine which SDD (vpathN) volumes you will use as Oracle8 database devices.
4. Partition the selected volumes using the HP-UX format utility. If SDD raw partitions are to be used by Oracle8 as database devices, be sure to leave disk cylinder 0 of the associated volume unused. This protects UNIX disk labels from corruption by Oracle8, as described in the *Oracle8 Installation Guide*.
5. Ensure that the Oracle software owner has read and write privileges to the selected SDD raw partition device files under the /devices directory.
6. Set up symbolic links from the oradata directory (under the first of the three mount points). Link the database files system <db>.dbf, tempdb.dbf, rbsdb.dbf, toolsd.bdbf, and usersdb.dbf to SDD raw device links (files named /dev/rdsk/vpathNs). Point to the partitions of the appropriate size, where *db* is the name of the database that you are creating. (The default is test.)
7. Install the Oracle8 server following the instructions in the *Oracle8 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 steps 2 and 6 for the redo logs and 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: Perform 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 "Installing SDD" on page 144.
3. Change to the directory where you installed the SDD utilities. Issue the **showvpath** command.
4. Check the output of the **showvpath** command to see whether you find a cntndn directory that is the same as the one where the Oracle files are.
5. Use the SDD partition identifiers instead of the original HP-UX identifiers when mounting the file systems.

   If you originally used the following HP-UX identifiers:

   mount /dev/dsk/c1t3d2 /oracle/mp1

   Replace those with the following SDD partition identifiers:

   mount /dev/dsk/vpath2 /oracle/mp1

   For example, assume that you found that vpath2 was the SDD identifier.

Follow the instructions in the *Oracle Installation Guide* for setting ownership and permissions.

If using raw partitions: Perform the following procedure if you have Oracle8 already installed and want to reconfigure it to use SDD partitions instead of sdisk partitions (for example, partitions accessed through /dev/rdsk/cntndn files).

All Oracle8 control, log, and data files are accessed either directly from mounted file systems or using links from the oradata subdirectory of each Oracle mount point that is set up on the server. Therefore, the process of converting an Oracle installation from sdisk to SDD has two parts:

- Change the Oracle physical devices for the mount points in /etc/fstab from sdisk device partition links to the SDD device partition links that access the same physical partitions.
- Re-create links to raw sdisk device links to point to raw SDD device links that access the same physical partitions.

Converting an Oracle installation from sdisk to SDD: Perform the following conversion steps:

1. Back up your Oracle8 database files, control files, and redo logs.

2. Obtain the sdisk device names for the Oracle8 mounted file systems by looking up the Oracle8 mount points in /etc/fstab and extracting the corresponding sdisk device link name (for example, /dev/rdsk/c1t4d0).
3. Launch the sqlplus utility.
4. Enter the **select * from sys.dba_data_files;** command.

   Determine the underlying device where each data file resides, either by looking up mounted file systems in /etc/fstab or by extracting raw device link names directly from the **select** command output.
5. Fill in the following table for planning purposes:

Oracle device link	File attributes			SDD device link
	Owner	Group	Permissions	
/dev/rdsk/c1tld0	oracle	dba	644	/dev/rdsk/vpath4

6. Fill in column 2 by issuing the command **ls -l** on each device link listed in column 1 and extracting the link source device file name.
7. Fill in the File Attributes columns by issuing the command **ls -l** on each Actual Device Node from column 2.
8. Install SDD following the instructions in the “Installing SDD” on page 144.
9. Fill in the SDD Device Links column by matching each **cntndnsn** device link listed in the Oracle Device Link column with its associated **vpathN** device link name by entering the **/opt/IBMsdd/bin/showvpath** command:
10. Fill in the SDD Device Nodes column by issuing the command **ls -l** on each SDD Device Link and tracing back to the link source file.
11. Change the attributes of each node listed in the SDD Device Nodes column to match the attributes listed to the left of it in the File Attributes column using the UNIX **chown, chgrp**, and **chmod** commands.
12. Make a copy of the existing /etc/fstab file. Edit the /etc/fstab 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 sdisk device link listed in the Oracle Device Link column.
14. Restart the server.
15. Verify that all file system and database consistency checks complete successfully.

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 Web site:

www-1.ibm.com/servers/storage/support/software/sdd.html

Verifying the hardware and software requirements

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

Hardware

- 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 fibre-channel adapters.

  Go to the following Web site for more information about the fibre-channel adapters that you can use on your Linux host system:

  www-1.ibm.com/servers/storage/support/software/sdd.html
- Subsystem LUNs that have been created and configured for multiport access. Each LUN can have up to 32 sd instances, with one for each path on the server.
- A fibre-optic cable to connect each fibre-channel adapter to an ESS port, or to switch ports subsequently zoned to ESS ports.

Software

- SuSE SLES7 with kernel k_smp-2.4.18-224 kernel rpm
- SLES8 / UnitedLinux 1.0 kernels on i686:
  - k_smp-2.4.21-138
  - k_smp-2.4.21-169
  - k_smp-2.4.21-190
  - k_smp-2.4.21-198
- pSeries kernels:
  - SLES8 kernel-ppc64-2.4.21-111
  - SLES8 kernel-ppc64-2.4.21-147
  - SLES8 kernel-ppc64-2.4.21-154
- 32-bit, Intel-based, Red Hat 7.x Linux operating system running:
  - kernel-2.4.18-24.7.xbigmem
  - kernel-2.4.18-27.7.xbigmem
- Red Hat Linux Advanced Server 2.1 with:
  - 2.4.9-e.30enterprise
  - 2.4.9-e.34enterprise
  - 2.4.9-e.35enterprise
  - 2.4.9-e.38enterprise
- Red Hat Linux Advanced Server 2.1 with:

© Copyright IBM Corp. 1999, 2004

  - 2.4.9-3.30summit
  - 2.4.9-3.34summit
  - 2.4.9-3.35summit
  - 2.4.9-e.38summit
- 32-bit, Intel-based, Red Hat EL3 running:
  - 2.4.21-9.ELsmp
  - 2.4.21-9.ELhugemem
  - 2.4.21-9.0.1.ELsmp
  - 2.4.21-9.0.1.ELhugemem
- Red Hat EL3 (ppc64) running:
  - 2.4.21-9.EL
  - 2.4.21-9.0.1.EL

Unsupported environments

SDD does not support the following environments:

- A host system with both a SCSI and fibre-channel connection to a shared ESS LUN
- Logical Volume Manager (LVM)
- The EXT3 file system on an SDD pseudo-device is **only** supported on:

  SLES8/UL1 SP3 or newer (kernel 2.4.21-138 or newer)
- Single-path mode during concurrent download and installation of LIC nor during any ESS concurrent maintenance that impacts the path attachment, such as an ESS host-bay-adapter replacement

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 the ESS

Before you install SDD, configure your ESS for multiport access for each LUN. SDD requires a minimum of two independent paths 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 ESS ports is considered a multipath fibre-channel connection.

Refer to the *IBM TotalStorage Enterprise Storage Server: Introduction and Planning Guide* for more information about how to configure the ESS.

Refer to the *IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide* for information on working around Linux LUN limitations.

Configuring the SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000

Before you install SDD, configure your SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 for multiport access for each LUN. SDD requires a

minimum of two independent paths 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 SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 ports is considered a multipath fibre-channel connection.

For information about configuring your SAN Volume Controller, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Configuration Guide*.

For information about configuring your SAN Volume Controller for Cisco MDS 9000, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Configuration Guide*.

Refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Host Systems Attachment Guide* for information on working around Linux LUN limitations.

Configuring the SAN File System

SAN File System metadata servers already have SDD pre-installed and configured. The SAN File System may have specific configuration and support requirements for its Linux Client systems. Refer to the publications in Table 3 on page xxiii for specific Linux host system requirements and for information about upgrading SDD on the SAN File System metadata servers.

Configuring fibre-channel adapters for ESS

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.

Refer to 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 for SAN Volume Controller

For information about configuring your SAN Volume Controller, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Planning Guide* and the *IBM TotalStorage Virtualization Family: SAN Volume Controller Configuration Guide*.

Refer to the *IBM TotalStorage Virtualization Family: 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.

Configuring fibre-channel adapters for SAN Volume Controller for Cisco MDS 9000

For information about configuring your SAN Volume Controller for Cisco MDS 9000, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Planning Guide* and the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Configuration Guide*.

Refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 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.

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.

Perform 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 compact disc (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: enter **cd /mnt/cdrom**
   - For SuSE: enter **cd /media/cdrom**
5. If you're running Red Hat, enter **cd redhat**; if you're 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***-1.i686.redhat.rpm** or **rpm -qpl IBMsdd-***N.N.N.N***-1.i686.suse.rpm** to view all the files in the package,

   where *N.N.N.N* represents the current version release modification level number; for example, *N.N.N.N* = 1.3.0.1.
7. Enter **rpm -iv IBMsdd-***N.N.N.N***-1.i686.redhat.rpm** or enter **rpm -iv IBMsdd-***N.N.N.N***-1.i686.suse.rpm** to install SDD.

   A message similar to the following is displayed:

```
Preparing for installation ...
IBMsdd-N.N.N.N-1
```

Upgrading the SDD

Perform 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: enter **cd /mnt/cdrom**
   - For SuSE: enter **cd /media/cdrom**
5. If you're running Red Hat, enter **cd redhat**; if you're 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***-1.i686.redhat.rpm** or **rpm -qpl IBMsdd-***N.N.N.N***-1.i686.suse.rpm** to view all the files in the package.

   (where *N.N.N.N* represents the current version release modification level number; *N.N.N.N* = 1.3.0.1 for example.)
7. Enter **rpm -U IBMsdd-***N.N.N.N***-1.i686.redhat.rpm** or enter **rpm -U IBMsdd-***N.N.N.N***-1.i686.suse.rpm** to upgrade SDD.

   A message similar to the following is displayed:

```
Preparing for installation ...
IBMsdd-N.N.N.N-1
```

Verifying the SDD installation

The SDD installation installs the device driver and its utilities in the /opt/IBMsdd directory. Table 26 lists the SDD driver and its major component files.

Table 26. SDD components for a Linux host system

File name	Location	Description
sdd-mod.o-xxxy	/opt/IBMsdd	SDD device driver file (where XXX stands for the kernel level of your host system and y represents smp or up
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

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        88190 Sep 26 17:40 sdd-mod.o-2.4.2
-rw-r-----    1 root     root        88817 Sep 26 17:40 sdd-mod.o-2.4.2-smp
-rw-r-----    1 root     root        88094 Sep 26 17:40 sdd-mod.o-2.4.6
-rw-r-----    1 root     root        88689 Sep 26 17:40 sdd-mod.o-2.4.6-smp
-rw-r-----    1 root     root        88775 Sep 26 17:40 sdd-mod.o-2.4.9
-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
```

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.

You can manually or automatically load and configure SDD on your host Linux system. Manual configuration requires that you use a set of SDD-specific commands while automatic configuration requires a system restart. This section provides instructions for the following procedures:

- Using multiple commands for SDD configuration
- Using a single command for SDD configuration
- Configuring SDD at system startup
- Maintaining vpath device configuration persistence

If you are loading and configuring SDD for the first time, you should manually configure SDD with a set of SDD commands. This manual configuration process enables you to become familiar with the useful SDD commands in Table 27.

Table 27. Summary of SDD commands for a Linux host system

Command	Description
cfgvpath	Configures vpath devices.[1]
cfgvpath query	Displays all sd devices.
lsvpcfg	Displays the current devices that are configured and their corresponding paths.
rmvpath	Removes one or all vpath devices.
addpaths	Adds any new paths to an existing vpath.
sdd start	Loads the SDD driver and automatically configures disk devices for multipath access.
sdd stop	Unloads the SDD driver (requires no vpath devices currently in use).
sdd restart	Unloads the SDD driver (requires no vpath devices currently in use), and then loads the SDD driver and automatically configures disk devices for multipath access.

Note: [1] The vpath devices are assigned names according to the following scheme:

vpatha, vpathb,...,vpathp, vpathaa, vpathab,...,vpathap, vpathba, vpathbb,...,vpathbp,...,

Using multiple commands for SDD configuration

Perform the following steps to manually load and configure SDD on your Linux host system:

1. Log on to your Linux host system as the root user.
2. Change to the /opt/IBMsdd directory and enter **insmod ./filename**, where *filename* is the name of the binary module with the same version as your kernel (sdd-mod.o-2.4.2-smp for example). If you are unsure of your kernel version, use the command **uname -a** to list it. The example output here shows a kernel version 2.4.18 that is compiled for smp:

```
Linux hostname 2.4.18 #1 SMP Tue Mar 5 11:02:41 PST 2002 i686 unknown
```

In this example, the command **insmod ./sdd-mod.o-2.4.18-smp** would load the correct SDD driver code.

3. 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)
```

4. 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 1.3.0.0 is installed on a Linux host system running a 2.4.9 uniprocessor kernel:

```
sdd-mod: SDD 1.3.0.0 2.4.9 UP  Sep 26 2001 17:39:06 (C) IBM Corp.
```

5. Enter **cat /proc/scsi/scsi** to verify that all expected devices (LUNs) were detected and configured.

   Enter **cat /proc/scsi/xxx/N** to display the status of a specific adapter and the names of the attached devices. `xxx` in the **cat /proc/scsi/xxx/*N*** command stands for the adapter driver name, while *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,

.
.
.
```

6. Enter **cfgvpath query** to verify that you have configured the SCSI disk devices that you allocated and configured for SDD.

   After you enter the **cfgvpath query** command, a message similar to the following is displayed. This example output is for a system with ESS, SAN Volume Controller, and SAN Volume Controller for Cisco MDS 9000 LUNs.

```
/dev/sda  (  8,   0) host=0  ch=0  id=0  lun=0   vid=IBM      pid=DDYS-T36950M serial=xxxxxxxxxxxx                    ctlr_flag=0 ctlr_nbr=0 df_ctlr=0 X
/dev/sdb  (  8,  16) host=2  ch=0  id=0  lun=0   vid=IBM      pid=2105E20      serial=60812028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdc  (  8,  32) host=2  ch=0  id=0  lun=1   vid=IBM      pid=2105E20      serial=70912028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdd  (  8,  48) host=2  ch=0  id=0  lun=2   vid=IBM      pid=2105E20      serial=31B12028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sde  (  8,  64) host=2  ch=0  id=0  lun=3   vid=IBM      pid=2105E20      serial=31C12028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdf  (  8,  80) host=2  ch=0  id=1  lun=0   vid=IBM      pid=2105E20      serial=60812028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdg  (  8,  96) host=2  ch=0  id=1  lun=1   vid=IBM      pid=2105E20      serial=70912028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdh  (  8, 112) host=2  ch=0  id=1  lun=2   vid=IBM      pid=2105E20      serial=31B12028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdi  (  8, 128) host=2  ch=0  id=1  lun=3   vid=IBM      pid=2105E20      serial=31C12028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdj  (  8, 144) host=2  ch=0  id=6  lun=0   vid=IBM      pid=2145         serial=600507680183000a800000000000000a ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdk  (  8, 160) host=2  ch=0  id=6  lun=1   vid=IBM      pid=2145         serial=600507680183000a800000000000000b ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdl  (  8, 176) host=2  ch=0  id=6  lun=2   vid=IBM      pid=2145         serial=600507680183000a800000000000000c ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdm  (  8, 192) host=2  ch=0  id=6  lun=3   vid=IBM      pid=2145         serial=600507680183000a800000000000000d ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdn  (  8, 208) host=2  ch=0  id=6  lun=4   vid=IBM      pid=2145         serial=600507680183000a800000000000000e ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdo  (  8, 224) host=2  ch=0  id=6  lun=5   vid=IBM      pid=2145         serial=600507680183000a800000000000000f ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdp  (  8, 240) host=2  ch=0  id=6  lun=6   vid=IBM      pid=2145         serial=600507680183000a8000000000000010 ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdq  ( 65,   0) host=2  ch=0  id=6  lun=7   vid=IBM      pid=2145         serial=600507680183000a8000000000000011 ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdr  ( 65,  16) host=2  ch=0  id=6  lun=8   vid=IBM      pid=2145         serial=600507680183000a8000000000000012 ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sds  ( 65,  32) host=2  ch=0  id=6  lun=9   vid=IBM      pid=2145         serial=600507680183000a8000000000000013 ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdt  ( 65,  48) host=2  ch=0  id=7  lun=0   vid=IBM      pid=2145         serial=600507680183000a800000000000000a ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdu  ( 65,  64) host=2  ch=0  id=7  lun=1   vid=IBM      pid=2145         serial=600507680183000a800000000000000b ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdv  ( 65,  80) host=2  ch=0  id=7  lun=2   vid=IBM      pid=2145         serial=600507680183000a800000000000000c ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdw  ( 65,  96) host=2  ch=0  id=7  lun=3   vid=IBM      pid=2145         serial=600507680183000a800000000000000d ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdx  ( 65, 112) host=2  ch=0  id=7  lun=4   vid=IBM      pid=2145         serial=600507680183000a800000000000000e ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdy  ( 65, 128) host=2  ch=0  id=7  lun=5   vid=IBM      pid=2145         serial=600507680183000a800000000000000f ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdz  ( 65, 144) host=2  ch=0  id=7  lun=6   vid=IBM      pid=2145         serial=600507680183000a8000000000000010 ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdaa ( 65, 160) host=2  ch=0  id=7  lun=7   vid=IBM      pid=2145         serial=600507680183000a8000000000000011 ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdab ( 65, 176) host=2  ch=0  id=7  lun=8   vid=IBM      pid=2145         serial=600507680183000a8000000000000012 ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdac ( 65, 192) host=2  ch=0  id=7  lun=9   vid=IBM      pid=2145         serial=600507680183000a8000000000000013 ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdad ( 65, 208) host=2  ch=0  id=10 lun=0   vid=IBM      pid=2145         serial=600507680183000a800000000000000a ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdae ( 65, 224) host=2  ch=0  id=10 lun=1   vid=IBM      pid=2145         serial=600507680183000a800000000000000b ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdaf ( 65, 240) host=2  ch=0  id=10 lun=2   vid=IBM      pid=2145         serial=600507680183000a800000000000000c ctlr_flag=1 ctlr_nbr=0 df_ctlr=0

⋮
```

```
/dev/sda  (  8,   0) host=0  ch=0  id=0  lun=0   vid=IBM      pid=DDYS-T36950M serial=xxxxxxxxxxxx                    ctlr_flag=0 ctlr_nbr=0 df_ctlr=0 X
/dev/sdb  (  8,  16) host=2  ch=0  id=0  lun=0   vid=IBM      pid=2105E20      serial=60812028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdc  (  8,  32) host=2  ch=0  id=0  lun=1   vid=IBM      pid=2105E20      serial=70912028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdd  (  8,  48) host=2  ch=0  id=0  lun=2   vid=IBM      pid=2105E20      serial=31B12028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sde  (  8,  64) host=2  ch=0  id=0  lun=3   vid=IBM      pid=2105E20      serial=31C12028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdf  (  8,  80) host=2  ch=0  id=1  lun=0   vid=IBM      pid=2105E20      serial=60812028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdg  (  8,  96) host=2  ch=0  id=1  lun=1   vid=IBM      pid=2105E20      serial=70912028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdh  (  8, 112) host=2  ch=0  id=1  lun=2   vid=IBM      pid=2105E20      serial=31B12028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdi  (  8, 128) host=2  ch=0  id=1  lun=3   vid=IBM      pid=2105E20      serial=31C12028                        ctlr_flag=0 ctlr_nbr=0 df_ctlr=0
/dev/sdj  (  8, 144) host=2  ch=0  id=6  lun=0   vid=IBM      pid=2145         serial=600507680183000a800000000000000a ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdk  (  8, 160) host=2  ch=0  id=6  lun=1   vid=IBM      pid=2145         serial=600507680183000a800000000000000b ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdl  (  8, 176) host=2  ch=0  id=6  lun=2   vid=IBM      pid=2145         serial=600507680183000a800000000000000c ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdm  (  8, 192) host=2  ch=0  id=6  lun=3   vid=IBM      pid=2145         serial=600507680183000a800000000000000d ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdn  (  8, 208) host=2  ch=0  id=6  lun=4   vid=IBM      pid=2145         serial=600507680183000a800000000000000e ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdo  (  8, 224) host=2  ch=0  id=6  lun=5   vid=IBM      pid=2145         serial=600507680183000a800000000000000f ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdp  (  8, 240) host=2  ch=0  id=6  lun=6   vid=IBM      pid=2145         serial=600507680183000a8000000000000010 ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdq  ( 65,   0) host=2  ch=0  id=6  lun=7   vid=IBM      pid=2145         serial=600507680183000a8000000000000011 ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdr  ( 65,  16) host=2  ch=0  id=6  lun=8   vid=IBM      pid=2145         serial=600507680183000a8000000000000012 ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sds  ( 65,  32) host=2  ch=0  id=6  lun=9   vid=IBM      pid=2145         serial=600507680183000a8000000000000013 ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdt  ( 65,  48) host=2  ch=0  id=7  lun=0   vid=IBM      pid=2145         serial=600507680183000a800000000000000a ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdu  ( 65,  64) host=2  ch=0  id=7  lun=1   vid=IBM      pid=2145         serial=600507680183000a800000000000000b ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdv  ( 65,  80) host=2  ch=0  id=7  lun=2   vid=IBM      pid=2145         serial=600507680183000a800000000000000c ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdw  ( 65,  96) host=2  ch=0  id=7  lun=3   vid=IBM      pid=2145         serial=600507680183000a800000000000000d ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdx  ( 65, 112) host=2  ch=0  id=7  lun=4   vid=IBM      pid=2145         serial=600507680183000a800000000000000e ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdy  ( 65, 128) host=2  ch=0  id=7  lun=5   vid=IBM      pid=2145         serial=600507680183000a800000000000000f ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdz  ( 65, 144) host=2  ch=0  id=7  lun=6   vid=IBM      pid=2145         serial=600507680183000a8000000000000010 ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdaa ( 65, 160) host=2  ch=0  id=7  lun=7   vid=IBM      pid=2145         serial=600507680183000a8000000000000011 ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdab ( 65, 176) host=2  ch=0  id=7  lun=8   vid=IBM      pid=2145         serial=600507680183000a8000000000000012 ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdac ( 65, 192) host=2  ch=0  id=7  lun=9   vid=IBM      pid=2145         serial=600507680183000a8000000000000013 ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdad ( 65, 208) host=2  ch=0  id=10 lun=0   vid=IBM      pid=2062         serial=600507680183000a800000000000000a ctlr_flag=1 ctlr_nbr=0 df_ctlr=0
/dev/sdae ( 65, 224) host=2  ch=0  id=10 lun=1   vid=IBM      pid=2062         serial=600507680183000a800000000000000b ctlr_flag=1 ctlr_nbr=1 df_ctlr=0
/dev/sdaf ( 65, 240) host=2  ch=0  id=10 lun=2   vid=IBM      pid=2062         serial=600507680183000a800000000000000c ctlr_flag=1 ctlr_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.

7. Enter **cfgvpath** to configure SDD vpath devices.

   The configuration information will be 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 *your_configuration_file_name*.cfg** command where *your_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 vpath devices that this file defines.

   You can remove an SDD vpath device by using the **rmvpath xxx** command, where `xxx` represents the name of the vpath device that is selected for removal.

Enter **cfgvpath ?** or **rmvpath ?** for more information about the **cfgvpath** or **rmvpath** command.

8. Verify the vpath 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 ESS and SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 LUNs.

```
sdd-mod: SDD 1.4.0.0 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/sdbr /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/sdby /dev/sdci /dev/sdcs
```

   See Chapter 12, “Using the datapath commands,” on page 283 for more information about the **datapath query device** command and all other SDD datapath commands.

Using a single command for SDD configuration

You can also manually load and configure SDD by issuing the **sdd start** command. Successful execution of the **sdd start** command performs all the tasks described in “Using multiple commands for SDD configuration” on page 176.

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.

Configuring SDD at system startup

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.

Perform the following steps to configure SDD at your Red Hat or SuSE SLES8 system startup:

1. Log on to your Linux host system as the root user.
2. Enter **chkconfig --level** *X* **sdd on** to enable run level *X* at startup (where *X* represents the system run level). Refer to Linux system documentation for information about **chkconfig**.
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**

Perform the following steps to configure SDD at your SuSE SLES 7 system startup:

1. Log on to your Linux host system as the root user.
2. Enter **insserv sdd**.

3. Restart your host system so that SDD is loaded and configured.

If necessary, you can disable the startup option by entering:

insserv -r sdd

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 may be available from the HBA driver vendor.

Maintaining vpath device configuration persistence

Use the **cfgvpath** command to configure vpath devices. For first time configuration, the configuration method finds all sd devices, then configures and assigns vpath devices accordingly. The configuration is saved in /etc/vpath.conf to maintain name persistence in subsequent driver loads and configurations.

After the vpath devices are configured, issue **lsvpcfg** or the **datapath query device** command to verify the configuration. See "datapath query device" on page 291 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
#pathh 70C20530
```

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 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 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 303 for more information about the **datapath set device policy** command.

Dynamically adding paths to SDD devices

You can issue the **addpaths** command to add paths to SDD devices. For SDD to discover new paths, the Linux kernel SCSI disk driver must already be aware of the path. For information on dynamically adding an ESS disk to a running Linux system, refer to the *IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide* for details.

Dynamic reconfiguration

Dynamic reconfiguration provides a way to automatically detect path configuration changes without requiring a reboot.

1. **cfgvpath** :

   This operation finds the current hardware configuration and compares it to the vpath configuration in memory and then identifies a list of differences. It then issues commands to update the vpath configuration in memory with the current hardware configuration. The commands that **cfgvpath** issues to the vpath driver are:

   a. Add vpath device.
   b. Remove a vpath device; this will fail if device is busy.
   c. Add path to vpath device.
   d. Remove path for a vpath device; this will fail deletion of path if device is busy, but will set path to DEAD and OFFLINE.

2. The **rmvpath** command removes one or more vpath devices.

```
rmvpath                       # Remove all vpaths

rmvpath vpath_name            # Remove one vpath at a time
                              # this will fail if device is busy
```

Removing SDD

You must unload the SDD driver before removing SDD. Perform 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 **lsmod** to list the modules and version of SDD that is loaded.
3. Enter **rmmod** *filename*, where *filename* is the name of the binary module with the same version as your kernel (sdd-mod.o-2.4.2-smp for example). If you are unsure of your kernel version, use the command **uname -a** to list it. The example output here shows a kernel version 2.4.18 that is compiled for smp:

   ```
   Linux hostname 2.4.18 #1 SMP Tue Mar 5 11:02:41 PST 2002 i686 unknown
   ```

   In this example, the command **rmmod ./sdd-mod.o-2.4.18-smp** would remove the correct SDD driver code.
4. Enter **rpm -e IBMsdd** to remove the SDD package.
5. Verify the SDD removal by entering either **rpm -q IBMsdd** or **rpm -ql IBMsdd**.

   If you successfully removed SDD, output similar to the following is displayed:

   ```
   package IBMsdd is not installed
   ```

Note: The `rmmod` command will not unload a driver that is in use.

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 11, “Using the SDD server and the SDDPCM server,” on page 279 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, then 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 183 to confirm that the SDD server started successfully.

Changing to a different port number for the SDD server

See "Changing the sddsrv or pcmsrv TCP/IP port number" on page 282.

Stopping the SDD server

To stop the SDD server:

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. Execute **telinit q**.

See "Verifying if the SDD server has started" on page 183 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**.

Understanding the SDD error recovery policy

SDD, when in *multipath mode*, makes it possible for you to concurrently download and install licensed internal code to the supported storage device while application I/O continues running. SDD makes this process transparent to the Linux host system through its error recovery algorithm.

The Linux host system has multiple paths that are configured to a LUN.

Important: I/O will be run to all available ESS paths.

Only paths to one of the nodes in a SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 pair will be used at any one time. If that node becomes unavailable, paths to the remaining node will be used.

SDD in *multipath* mode has the following characteristics:

Both the Linux SDD 1.3.1.0 (or later) and Linux SDD 1.3.0.1 (or earlier) error recovery policies in multiple-path mode have the following common characteristics:

- If an I/O error occurs on the last operational path to a device, SDD attempts to reuse (or perform a failback operation to return to) a previously-failed path.

The Linux SDD 1.3.1.0 (or later) error recovery policy in *multiple-path mode* has the following latest characteristics:

- If an I/O error occurs on a path, Linux SDD 1.3.1.0 (or later) does not attempt to use the path until three successful I/O operations on an operational path.
- If an I/O error occurs consecutively on a path and the I/O error count reaches three, SDD immediately changes the state of the failing path to DEAD.

- Both the SDD driver and the SDD server daemon can put a last path into DEAD state, if this path is no longer functional. The SDD server can automatically change the state of this path to OPEN after it is recovered. Alternatively, you can manually change the state of the path to OPEN after it is recovered by using **datapath set path online** command. Go to “datapath set device path” on page 304 for more information.
- If the SDD server daemon detects that the last CLOSE path is failing, the daemon will change the state of this path to CLOSE_DEAD. The SDD server can automatically recover the path if it is detected that it is functional.
- If an I/O fails on all OPEN paths to a LUN, SDD returns the failed I/O to the application and changes the state of all OPEN paths (for failed I/Os) to DEAD, even if some paths did not reach an I/O error count of three.
- If an OPEN path already failed some I/Os, it will not be selected as a retry path.

The Linux SDD 1.3.0.1 (or earlier) error recovery policy in *multiple-path mode* has the following characteristics:

- If an I/O error occurs on a path, Linux SDD 1.3.0.1 does not attempt to use the path until 2000 successful I/O operations on an operational path.
- The last path is reserved in OPEN state.
- If an I/O fails on all OPEN paths to a LUN, SDD returns the failed I/O to the application and leaves all the paths in OPEN state.
- A failed I/O is retried on all OPEN paths to a LUN even if the OPEN path already failed I/Os.
- SDD changes the failed path from the DEAD state back to the OPEN state after 50 000 successful I/O operations on an operational path.

Understanding the SDD support for single-path configuration

SDD does not support concurrent download and installation of the licensed internal code (LIC) on single-path mode.

However, SDD supports *single-path* SCSI or fibre-channel connection from your Linux host system to an ESS and *single-path* fibre-channel connection from your Linux host system to a SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000.

Notes:

1. SDD supports one SCSI adapter on the host system.
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.
3. Because of single-path connection, SDD can not provide single-point-failure protection and load balancing. IBM does not recommend this.

Partitioning SDD vpath devices

Disk partitions are known as *logical devices*. SDD for Linux allows configuration of whole devices only. 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/vpath*X* is created for each supported multipath SCSI disk device (where *X* represents the unique device name; as with sd devices, *X* may be one or two letters).

Device files /dev/vpath*XY* 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/vpath*XY* (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 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, vpath device names can replace corresponding sd 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.

Chapter 6. Using SDD on a NetWare host system

Attention: SDD does not support NetWare host systems attached to:
- SAN Volume Controller
- SAN Volume Controller for Cisco MDS 9000

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 an ESS. 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 SDD Web site at:

www-1.ibm.com/servers/storage/support/software/sdd.html

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 Web site:

www.ibm.com/servers/storage/support/

Hardware requirements

- IBM TotalStorage SAN Fibre Channel Switch 2109 is recommended
- Host system
- Fibre-channel switch
- SCSI adapters and cables
- Fibre-channel adapters and cables

Software requirements

- Windows operating system running on the client
- One of the following NetWare operating systems running on the server:
  - Novell NetWare 5.1 SP6
  - Novell NetWare 6 SP1 or SP2 or SP3
  - NetWare 6.5 SP1.1
- NetWare Cluster Service for NetWare 5.1 if servers are being clustered
- NetWare Cluster Service for NetWare 6.0 if servers are being clustered
- ConsoleOne
- SCSI and fibre-channel device drivers

© Copyright IBM Corp. 1999, 2004

Supported environments

SDD supports:

- Novell NetWare 5.1 SP6
- Novell NetWare 6 SP1 or SP2 or SP3
- Novell NetWare 6.5 SP1.1
- 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.

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 ESS LUN
- Single-path mode during concurrent download and installation of LIC nor during any ESS concurrent maintenance that impacts the path attachment, such as an ESS host-bay-adapter replacement

ESS requirements

- To successfully install SDD, ensure that the ESS devices are configured as either an:
  - IBM 2105xxx (SCSI-attached device)
  - IBM FC 2105xxx (fibre-channel attached device)

  Note: *xxx* represents the ESS 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 Web site:

www-1.ibm.com/servers/storage/support/software/sdd.html

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 Web site:

www.ibm.com/servers/storage/support/

Note: If your host has only one fibre-channel adapter, you need to connect through a switch to multiple ESS 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 an ESS 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 Web site:

www-1.ibm.com/servers/storage/support/software/sdd.html

Preparing for SDD installation

Before you install SDD, you must configure the ESS for your host system and attach required fibre-channel adapters.

Configuring the ESS

Before you install SDD, you must configure:

- The ESS to your host system and the required fibre-channel or SCSI adapters that are attached.
- The ESS 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.**

Refer to the *IBM TotalStorage Enterprise Storage Server: Introduction and Planning Guide* for more information about how to configure the ESS.

Refer to the *IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide* 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.

Refer to 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.

Refer to 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
```

Refer to 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.

Refer to 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 abends to result.

To ensure that SDD failover functions as designed and to prevent potential volume dismounts, hangs, or abends, 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

The installation CD 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.

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.

Configuring SDD

To load the SDD module, enter **load SDD**.

To unload the SDD module, enter **unload SDD**.

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
- ESS 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 12, “Using the datapath commands,” on page 283 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 ESS and, if so, adds it to the path group.

See Chapter 12, “Using the datapath commands,” on page 283

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.

ESS logical unit detection

The SDD works only with ESS logical units. The 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 ESS.

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 ESS 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 ESS-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 SDD

To remove 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 Novel 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 occured on device 0x3A path 2
(Adapter Error Code: 0x8007, Device Error Code: 0x0000)
[No sense data]
This path is in CLOSE state.
Error has occured 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 "TEMPPOOL"...
Aug 8, 2003 3:05:06 am NSS<COMN>-3.02-xxxx: comnPool.c[2516]
Pool TEMPPOOL: 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=poolname.
END:nss /poolactivate=temppool
Activating pool "TEMPPOOL"...
** 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 TEMPPOOL 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 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 SDD Web site:

www-1.ibm.com/servers/storage/support/software/sdd.html

Verifying the hardware and software requirements

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

Hardware

- 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 ESS 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 an ESS controller port or a fibre-channel switch connected with ESS or SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 port.

  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 host adapters if you are attaching to an ESS.

  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 SAN Volume Controller or a SAN Volume Controller for Cisco MDS 9000.

  SDD requires enabling the host-adapter-dynamic-binding feature to have the same system device names for the same LUNs.

Software

SDD supports:

- ESS on a SPARC system running Solaris 2.6, Solaris 7, Solaris 8, or Solaris 9

  Note: ESS is supported on Solaris 9 by SDD 1.5.0.4 or later.
- SAN Volume Controller on a SPARC system running 64-bit Solaris 8 or 64-bit Solaris 9

  Note: SAN Volume Controller is supported on Solaris 9 by SDD 1.5.1.2 or later.
- SAN Volume Controller for Cisco MDS 9000 on a SPARC system running 64-bit Solaris 8

Supported environments

SDD supports the following environments:

- 32-bit applications on Solaris 2.6

© Copyright IBM Corp. 1999, 2004

- 32-bit and 64-bit applications on Solaris 7, Solaris 8, and Solaris 9

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 and installation of LIC nor during any ESS concurrent maintenance that impacts the path attachment, such as an ESS host-bay-adapter replacement
- Single-path configuration for Fibre Channel

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 1.

Preparing for SDD installation

Before you install SDD, you must configure the ESS device or SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 to your host system and attach the required fibre-channel adapters.

Configuring the ESS

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 how to configure your ESS, refer to the *IBM TotalStorage Enterprise Storage Server: Introduction and Planning Guide*.

Configuring the SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000

Before you install SDD, configure your SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 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, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Configuration Guide*.

For information about configuring your SAN Volume Controller for Cisco MDS 9000, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Configuration Guide*.

Determining if the SDD server for Expert is installed

If you previously installed the SDD server (the standalone version) for IBM TotalStorage Expert V2R1 (ESS Expert) on your Solaris host system, you must remove this standalone 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 standalone version of the SDD server (for ESS Expert).

To determine if the standalone version of the SDD server is installed on your host system, enter:

pkginfo -i SDDsrv

If you previously installed the standalone 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 standalone version of the SDD server (for ESS Expert) is SDDsrvSUN*bb_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 standalone 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 standalone version of the SDD server (for ESS Expert) from your Solaris host system, see the IBM(R) SUBSYSTEM DEVICE DRIVER SERVER 1.0.0.0 (sddsrv) README for IBM TotalStorage Expert V2R1 at the following Web site:

  www-1.ibm.com/servers/storage/support/software/swexpert.html

For more information about the SDD server daemon, go to “SDD server daemon” on page 216.

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 three different ways. The way that you choose depends on the kind of software that you have installed:

Table 28 further describes the various installation scenarios and how you should proceed.

Table 28. SDD installation scenarios

Installation scenario	Description	How to proceed

Table 28. SDD installation scenarios (continued)

Scenario 1	• 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: 1. "Installing SDD" 2. "Standard UNIX applications" on page 217
Scenario 2	• SDD is not installed. • An existing volume manager, software application, or DBMS is installed that communicates directly to the sd interface.	Go to: 1. "Installing SDD" 2. "Using applications with SDD" on page 216
Scenario 3	SDD is installed.	Go to "Upgrading the SDD" on page 210.

Table 29 lists the installation package file names that come with SDD.

Table 29. Operating systems and SDD package file names

Operating system	Package file name
Solaris 2.6 or 32-bit Solaris 7/8/9	sun32bit/IBMsdd
64-bit Solaris 7/8	sun64bit/IBMsdd

For SDD to operate properly, ensure that the Solaris patches are installed on your operating system. Go to the following Web site for the latest information about Solaris patches:

http://sunsolve.sun.com

For more information on the Solaris patches, refer to the *IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide* or the *IBM TotalStorage Virtualization Family: 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.

Note: Note that SDD package name has changed from IBMdpo to IBMsdd for SDD 1.4.0.0 or later.

Perform the following steps to install SDD on your Solaris host system:

Note: If the OS is Solaris 2.6, the OS bit-level is 32-bit.

If the OS is Solaris 7/8/9, you can check the OS bit-level that is executing by issuing **# isainfo -kv**.

1. Make sure that the SDD CD is available.
2. Insert the CD 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 example,

   ```
   pkgadd -d /cdrom/cdrom0/sun32bit IBMsdd or
   pkgadd -d /cdrom/cdrom0/sun64bit 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 will be executed 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/pathtest
/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/pathtest
/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. Enter **q** and press **Enter** to proceed. A message similar to the following message is displayed:

```
*** IMPORTANT NOTICE ***
This machine must now be rebooted in order to ensure
sane operation. Execute
       shutdown -y -i6 -g0
and wait for the "Console Login:" prompt.

SDD is now installed. Proceed to post-installation.
```

Verifying the SDD installation

To verify the SDD installation, perform 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 executes, SDD is installed.

Uninstalling 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.

Perform 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 packages are available:
1 IBMcli ibm2105cli
         (sparc) 1.1.0.0
2 IBMsdd IBM SDD driver Version: May-10-2000 16:51
         (sparc) 1
```

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 will be executed 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.

Postinstallation

After you complete the installation, 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 must restart your system to ensure proper operation. Enter the command:

 # shutdown -i6 -g0 -y

SDD devices are found in the /dev/rdsk and /dev/dsk directories. The device is named according to the SDD instance number. A device with an instance number of 0 would be: /dev/rdsk/vpath0a where *a* denotes the slice. Therefore, /dev/rdsk/vpath0c would be instance zero 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. Those components and files are listed in the following tables.

Table 30. 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 30. SDD components installed for Solaris host systems (continued)

S20vpath-config	/etc/rcS.d	Boot initialization script **Note:** This script must come before other LVM initialization scripts.
sddsrv	/sbin/sddsrv	SDD server daemon
sample_sddsrv.conf	/etc/sample_sddsrv.conf	Sample SDD server config file

Table 31. 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 32. SDD commands and their descriptions for Solaris host systems

Command	Description
cfgvpath [-c]	Configures vpath devices using the following process: 1. Scan the host system to find all devices (LUNs) that are accessible by the Solaris host. 2. Determine which devices (LUNs) are the *same* devices that are accessible through different paths. 3. Create configuration file /etc/vpath.cfg to save the information about devices. With -c 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 -c option: **cfgvpath** initializes the SDD device driver vpathdd with the information stored in /etc/vpath.cfg and creates pseudo-vpath devices /devices/pseudo/vpathdd* **Note: cfgvpath** without -c 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	Reconfigures vpath devices if vpath devices exist. See “Option 2: Dynamic reconfiguration” on page 211. If no vpath devices exist, use **cfgvpath** without -r option.

Table 32. SDD commands and their descriptions for Solaris host systems (continued)

Command	Description
showvpath	Lists all SDD devices and their underlying disks.
vpathmkdev	Creates files vpathMsN in the /dev/dsk/ and /dev/rdsk/ directories by creating links to the pseudo-vpath devices /devices/pseudo/vpathdd* that are created by the SDD driver. Files vpathMsN in the /dev/dsk/ and /dev/rdsk/ directories provide block and character access to an application the same way as the cxtydzsn devices created by the system. **vpathmkdev** is executed automatically during SDD package installation and should be executed manually to update files vpathMsN after hardware reconfiguration.
datapath	SDD driver console command tool.
rmvpath [all, vpathname]	Removes vpath devices from the configuration. See “Option 2: Dynamic reconfiguration” on page 211.

If you are not using a volume manager, software application, or DBMS that communicates directly to the sd interface, then 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 216 and read the information specific to the application that you are using.

Upgrading the SDD

To upgrade SDD:

1. Stop I/O activity on all SDD devices.
2. Uninstall SDD using the procedure in “Uninstalling SDD” on page 207.
3. Install SDD using the procedure in “Installing SDD” on page 204.
4. Reboot the system after SDD is installed.

Configuring SDD

Before you start the SDD configuration process, make sure that you have successfully configured the ESS or SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 to which your host system is attached and that the ESS or SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 is operational.

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.

Option 1: Reconfigure the system and reconfigure SDD

Perform 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. Run the SDD utility to reconfigure SDD. Enter **cfgvpath -c** and press **Enter**.
4. Shut down the system. Enter **shutdown -i6 -g0 -y** and press **Enter**.
5. After the restart, change to the /opt/IBMsdd/bin directory by entering:

   ```
   cd /opt/IBMsdd/bin
   ```

6. For Solaris 7/8/9:

   a. Enter **devfsadm** and press **Enter** to reconfigure all the drives.

   For Solaris 6:

   a. Enter **drvconfig** and press **Enter**.

   b. Enter **devlinks** and press **Enter** to reconfigure all the drives.
7. Enter **vpathmkdev** and press **Enter** to create all the 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 SDD.

Tip: Before executing the following SDD dynamic reconfiguration commands, execute the **showvpath** and **datapath query device** commands and 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. **cfgvpath -r**

   Note: If there are no existing vpath devices, the **cfgvpath -r** command will not dynamically reconfigure new vpath devices. You should execute **cfgvpath** to configure new vpath devices. Then execute **devfsadm** (or **drvconfig** on Solaris 2.6) and **vpathmkdev**.

   This operation finds the current hardware configuration and compares it to the vpath configuration in memory and then works out a list of differences. It then issues commands to put the vpath configuration in memory up-to-date with the current hardware configuration. The **cfgvpath -r** operation issues these commands to the vpath driver:

   a. Add vpath device.

      If you are adding new vpath devices, you need to execute **devfsadm** (**drvconfig** on Solaris 2.6) and **vpathmkdev**.

   b. Remove a vpath device; this will fail if device is busy.

   c. Add path to vpath device.

      If the vpath device changes from single path to multiple paths, the path selection policy of the vpath device will be changed to load-balancing policy.

d. Remove path for a vpath device; this deletion of path will fail if device is busy, but will set path to DEAD and OFFLINE.

   Removing paths of a vpath device or removing a 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 vpath device removal failure, all the paths of the 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 vpath devices. When the system is rebooted, the new SDD configuration would be used to configure vpath devices.

2. **rmvpath** command removes one or more vpath devices.
   a. To remove all vpath devices that are not busy:

      ```
      # rmvpath -all
      ```

   b. To remove one vpath device if the vpath device is not busy:

      ```
      # rmvpath vpathname
      ```

      For example, **rmvpath vpath10** will remove vpath10.

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.

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 vpath device, to show the current active policy for that device. The output should look similar to this:

```
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 303 for more information about the **datapath set device policy** command.

Preferred node path-selection algorithm for the SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000

The SAN Volume Controller and the SAN Volume Controller for Cisco MDS 9000 is a two-controller disk subsystem. SDD distinguishes the paths to a SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 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, then 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.

Understanding the SDD 1.3.1.0 (or later) support for single-path configuration for ESS

SDD 1.3.1.0 (or later) does not support concurrent download and installation of the ESS licensed internal code (LIC) on single-path mode. However, SDD supports single-path SCSI or fibre-channel connection from your Solaris host system to an ESS. It is possible to create a volume group or a vpath device with only a single path.

Notes:

1. SDD supports one SCSI adapter on the host system.
2. If your host has only one fibre-channel adapter, it requires you to connect through a switch to multiple ESS ports. You should have at least two fibre-channel adapters to prevent data loss due to adapter hardware failure or software failure. 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.**

Understanding the SDD error recovery policy

There are differences in the way that SDD 1.3.2.1 (or later) and SDD 1.3.1.7 (or earlier) handles error recovery for host systems:

SDD 1.3.1.7 (or earlier) error recovery policy

With SDD 1.3.1.7 (or earlier), error recovery policy is designed to cover a transient kind of error from the user applications. The error recovery policy prevents a path from becoming disabled in an event of transient errors.

Note: This policy can cause longer interruption for recovery to take place. The recovery policy halts the I/O activities on functional paths and vpath devices for some time before failing paths are set to DEAD state.

SDD 1.3.2.1 (or later) error recovery policy

With SDD 1.3.2.1 (or later), error recovery policy is designed to report errors to applications more quickly. With SDD 1.3.2.1 (or later), the applications receive failed I/O requests more quickly. This process prevents unnecessary retries, which can cause the I/O activities on good paths and vpath devices to halt for an unacceptable period of time.

Both the SDD 1.3.2.1 (or later) and SDD 1.3.1.7 (or earlier) error recovery policies support the following modes of operation:

single-path mode (for ESS only)

Note: SDD no longer supports single-path concurrent download of licensed internal code.

A Solaris host system has only one path that is configured to an ESS logical unit number (LUN).

In single-path mode, SDD has the following characteristics:

- When an I/O error occurs, SDD retries the I/O operation up to two times.
- With the SDD 1.3.2.1 (or later) error recovery policy, SDD returns the failed I/O to the application and sets the state of this failing path to DEAD. SDD relies on the SDD server daemon to detect the recovery of the single path. The SDD server daemon recovers this failing path and changes its state to OPEN. (SDD can change a single and failing path into DEAD state.)
- With the SDD 1.3.1.7 (or earlier) error recovery policy, SDD returns the failed I/O to the application and leaves this path in OPEN state (SDD never puts this single path into DEAD state).
- With SDD 1.3.2.1 (or later) the SDD server daemon detects the single CLOSE path that is failing and changes the state of this failing path to CLOSE_DEAD. When the SDD server daemon detects a CLOSE_DEAD path recovered, it will change the state of this path to CLOSE. With a

single path configured, the SDD vpath device can not be opened if it is the only path in a CLOSE_DEAD state.

multipath mode

The host system has multiple paths that are configured to an ESS LUN or to a SAN Volume Controller LUN or to a SAN Volume Controller for Cisco MDS 9000 LUN.

Both the SDD 1.3.2.1 (or later) and SDD 1.3.1.7 (or earlier) error recovery policies in multiple-path mode have the following common characteristics:

- If an I/O error occurs on the last operational path to a device, SDD attempts to reuse (perform a failback operation to return to) a previously-failed path.

The SDD 1.3.2.1 (or later) error recovery policy in multipath mode has the following latest characteristics:

- If an I/O error occurs on a path, SDD 1.3.2.1 (or later) does not attempt to use the path until three successful I/O operations occur on an operational path.
- If an I/O error occurs consecutively on a path and the I/O error count reaches three, SDD immediately changes the state of the failing path to DEAD.
- Both SDD driver and the SDD server daemon can put a last path into DEAD state, if this path is no longer functional. The SDD server can automatically change the state of this path to OPEN after it is recovered. Alternatively, you can manually change the state of the path to OPEN after it is recovered by using **datapath set path online** command. Go to "datapath set device path" on page 304 for more information.
- If the SDD server daemon detects that the last CLOSE path is failing, the daemon will change the state of this path to CLOSE_DEAD. The SDD server can automatically recover the path if it is detected that it is functional.
- If an I/O fails on all OPEN paths to an ESS LUN, SDD returns the failed I/O to the application and changes the state of all OPEN paths (for failed I/Os) to DEAD, even if some paths did not reach an I/O error count of three.
- If an OPEN path already failed some I/Os, it will not be selected as a retry path.

The SDD 1.3.1.7 (or earlier) error recovery policy in multipath mode has the following characteristics:

- If an I/O error occurs on a path, SDD 1.3.1.x does not attempt to use the path until 2000 successful I/O operations on an operational path.
- The last path is reserved in OPEN state.
- If an I/O fails on all OPEN paths to an ESS LUN, SDD returns the failed I/O to the application and leaves all the paths in OPEN state.
- A failed I/O is retried on all OPEN paths to an ESS LUN even if the OPEN path already failed I/Os.
- SDD changes the failed path from the DEAD state back to the OPEN state after 50 000 (200 000 for SDD 1.3.0.x or earlier) successful I/O operations on an operational path.

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 11, “Using the SDD server and the SDDPCM server,” on page 279 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 l 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.”

Starting the SDD server manually

If the SDD server does not start automatically after you perform 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. Execute **init q**.
4. Follow the directions in “Verifying if the SDD server has started” to confirm that the SDD server started successfully.

Changing to a different port number for the SDD server

See “Changing the sddsrv or pcmsrv TCP/IP port number” on page 282.

Stopping the SDD server

Perform the following steps to stop the SDD server :

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. Execute **init q**.
4. Check if sddsrv is running by executing **ps -ef lgrep sddsrv**. If sddsrv is still running, execute **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
- NFSs
- Oracle

Standard UNIX applications

If you have not already done so, install SDD using the procedure in “Installing SDD” on page 204. 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 SDD on a Network File System file server

The procedures in this section show how to install SDD for use with an exported file system (Network File System file server).

Setting up the Network File System for the first time

Perform 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 204 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/rdsk/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

Perform the following steps if you have the Network File System file server already configured to:

- Export file systems that reside on a multiport subsystem
- 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, you should have a copy of the *Veritas Volume Manager System Administrator's Guide* and the *Veritas Volume Manager Command Line Interface for Solaris*. These publications can be found at the following Web site:

www.veritas.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 204, 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. SVC 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 204.
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 super-user privileges to perform 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 differs depending on the choice you make.

Using a file system:

1. If you have not already done so, install SDD using the procedure described in “Installing SDD” on page 204.
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 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 204 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/rdsk/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 on page 219. Link the database files to SDD raw device links (files named /dev/rdsk/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: Perform 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 204 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: Perform 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/rdsk/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: Perform 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/rdsk/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/rdsk/c1t1d0s4

A message similar to the following message is displayed:

```
/dev/rdsk/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/rdsk/c1t1d0s4

   A message similar to the following message is displayed:

   ```
   crw-r--r-- oracle dba 32,252 Nov 16 11:49 /dev/rdsk/c1t1d0s4
   ```

7. Complete the basic installation steps in the “Installing SDD” on page 204 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 cntndnsn 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 will apply 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 super-user privileges to perform these procedures.

Note: SDD only supports Solstice DiskSuite line command interface. The DiskSuite Tool (metatool) does not recognize and present SDD devices for configuration.

Installing Solaris Volume Manager for the first time

Perform 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 204 section, if you have not already done so.
2. Configure the SPARC server to recognize all devices over all paths using the **boot -r** command.
3. Install the Solaris Volume Manager packages and the answerbook. Do not restart yet.
4. Determine which vpath devices 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 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. Perform the following instructions for setting up replica databases on the vpathNs partitions, where *N* is the 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 a vpath device correspond to slices 0 - 7 of the underlying multiport subsystem device.
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 used by DiskSuite into the md.tab file.

Installing SDD on a system that already has Solstice DiskSuite in place

Perform 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 204 section, if you have not already done so. After the installation completes, enter **shutdown -i6 -y -g0** and press **Enter**. This verifies the 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 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 vpath in case of 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 metadevice:

   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 perform these procedures.

Perform the following steps if you are installing a new UFS logging file system on vpath devices:

1. Install SDD using the procedure in the “Installing SDD” on page 204 section, if you have not already done so.
2. Determine which vpath (vpathNs) volumes that you will use as file system devices. Partition the selected 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 Solairs 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 <metadevice 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

Perform 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 204 section, if you have not already done so.
4. Match the sd device link names found in step 2 with vpath device link names (files named /dev/(r)dsk/vpathNs) by executing 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 corresponding 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 SDD on a Windows NT host system

Attention:	SDD does not support Windows NT clustering for systems attached to the SAN Volume Controller for Cisco MDS 9000

This chapter provides procedures for you to install, configure, remove, and use the SDD on a Windows NT host system that is attached to an ESS device or a SAN Volume Controller for Cisco MDS 9000. For updated and additional information that is not included in this chapter, see the Readme file on the CD-ROM or visit the SDD Web site at

www-1.ibm.com/servers/storage/support/software/sdd.html

Click **Subsystem Device Driver**.

Verifying the hardware and software requirements

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

Hardware

- One or more supported storage devices
- Host system
- For ESS devices: SCSI adapters and cables
- Fibre-channel adapters and cables

Software

- Windows NT 4.0 operating system with Service Pack 6A or later
- For ESS devices: SCSI device drivers
- Fibre-channel device drivers

Unsupported environments

SDD does not support the following environments:

- A host system with both a SCSI channel and fibre-channel connection to a shared LUN.
- SDD does not support I/O load balancing in a Windows NT clustering environment.
- You cannot store the Windows NT operating system or a paging file on an SDD-controlled multipath device (that is, SDD does not support boot from ESS device).
- Single-path mode during concurrent download and installation of LIC nor during any ESS concurrent maintenance that impacts the path attachment, such as an ESS host-bay-adapter replacement.
- Clustering on SAN Volume Controller for Cisco MDS 9000 devices

© Copyright IBM Corp. 1999, 2004

ESS requirements

To successfully install SDD, ensure that your host system is configured to the ESS as an Intel-based PC (personal computer) server with Windows NT 4.0 Service Pack 6A (or higher) installed.

Host system requirements

To successfully install SDD, your Windows NT host system must be an Intel-based system with Windows NT Version 4.0 Service Pack 6A or higher installed.

To install all components, you must have 1 MB (MB equals approximately 1 000 000 bytes) of disk space available.

The host system can be a uniprocessor or a multiprocessor system.

SCSI requirements

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

- No more than 32 SCSI adapters are attached.
- 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.

  Note: SDD also supports one SCSI adapter on the host system. With single-path access, concurrent download of licensed internal code is supported with SCSI devices. However, the load-balancing and failover features are not available.
- For information about the SCSI adapters that can attach to your Windows NT host system, go to the following Web site:

  www.ibm.com/storage/hardsoft/products/ess/supserver.htm

Fibre-channel requirements

To use the SDD 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 an supported storage device port.
- If you need the SDD I/O load-balancing and failover features, ensure that a minimum of two fibre-channel paths are configured between the host and the subsystem.

Note: If your host has only one fibre-channel adapter, it requires you to connect through a switch to multiple supported storage device ports. 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 the fibre-channel adapters that can attach to your Windows NT host system, go to the following Web site:

www.ibm.com/storage/hardsoft/products/ess/supserver.htm

Preparing for SDD installation

Before you install SDD, you must configure the supported storage device to your host system and the required fibre-channel adapters that are attached.

Configuring the ESS

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 LUN to use the load-balancing and failover features. **With a single path, failover protection is not provided.**

For information about configuring your ESS, refer to the *IBM TotalStorage Enterprise Storage Server: Introduction and Planning Guide*.

Configuring the SAN Volume Controller for Cisco MDS 9000

For information about configuring your SAN Volume Controller for Cisco MDS 9000, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Configuration Guide*.

Configuring fibre-channel adapters

You must configure the fibre-channel adapters that are attached to your Windows NT host system before you install SDD. Follow the adapter-specific configuration instructions to configure the adapters attached to your Windows NT host systems.

SDD supports Emulex adapters with full port driver only. When you configure the Emulex adapter for multipath functions, select **Allow Multiple paths to SCSI Targets** in the Emulex Configuration Tool panel.

Make sure that your Windows NT host system has Service Pack 6A or higher. Refer to the *IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide* for more information about installing and configuring fibre-channel adapters for your Windows NT host systems.

For information about installing and configuring fibre-channel adapters for your Windows NT host systems for the SAN Volume Controller for Cisco MDS 9000, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Configuration Guide*.

Configuring SCSI adapters for ESS devices

Attention: Failure to disable the BIOS of attached nonstart devices may cause your system to attempt to start from an unexpected nonstart device.

Before you install and use SDD, you must configure your SCSI adapters. For SCSI adapters that attach start devices, ensure that the BIOS for the adapter is *enabled*. For all other adapters that attach nonstart devices, ensure that the BIOS for the adapter is *disabled*.

Note: When the adapter shares the SCSI bus with other adapters, the BIOS must be disabled.

Installing SDD

The following section describes how to install SDD. Make sure that all hardware and software requirements are met before you install the Subsystem Device Driver. See “Verifying the hardware and software requirements” on page 245 for more information.

Perform the following steps to install the SDD filter and application programs on your system:

1. Log on as the administrator user.
2. Insert the SDD installation compact disc into the CD-ROM drive.
3. Start the Windows NT Explorer program.
4. Double-click the CD-ROM drive. A list of all the installed directories on the compact disc is displayed.
5. Double-click the \winNt\IBMsdd directory.
6. Run the setup.exe program. The Installshield program starts.
7. Click **Next**. The Software License agreement is displayed.
8. Select **I accept the terms in the License Agreement** and then click **Next**. The User Information window opens.
9. Type your name and your company name.
10. Click **Next**. The Choose Destination Location window opens.
11. Click **Next**. The Setup Type window opens.
12. Select the type of setup that you prefer from the following setup choices. IBM recommends that you select **Complete**.
    Complete
        Selects all options.
    Custom
        Select the options that you need.
13. Click **Next**. The Ready to Install The Program window opens.
14. Click **Install**. The Installshield Wizard Completed window opens.
15. Click **Finish**. The Installation program prompts you to restart your computer.
16. Click **Yes** to start your computer again. When you log on again, you see a **Subsystem Device Driver Management** entry in your Program menu containing the following files:
    a. Subsystem Device Driver Management
    b. Subsystem Device Driver manual
    c. README

Note: You can use the **datapath query device** command to verify the SDD installation. SDD is successfully installed if the command runs successfully.

Configuring SDD

To activate SDD, you need to restart your Windows NT system after it is installed. In fact, a restart is required to activate multipath support whenever a new file system or partition is added.

Adding paths to SDD devices

Attention: Ensure that SDD is installed *before* you add a new path to a device. Otherwise, the Windows NT server could lose the ability to access existing data on that device.

This section contains the procedures for adding paths to SDD devices in multipath environments.

Reviewing the existing SDD configuration information

Before adding any additional hardware, review the configuration information for the adapters and devices currently on your Windows NT server.

Verify that the number of adapters and the number of paths to each supported storage device volume match the known configuration. Perform the following steps to display information about the adapters and devices:

1. Click **Start → Program → Subsystem Device Driver → Subsystem Device Driver Management**. An MS-DOS window opens.
2. Enter **datapath query adapter** and press **Enter**. The output includes information about all the installed adapters. In the example shown in the following output, one host bus adapter has 10 active paths:

```
Active Adapters :1

Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active
    0  Scsi Port6 Bus0  NORMAL   ACTIVE        542          0     10      10
```

3. Enter **datapath query device** and press **Enter**. In the example shown on page 231, SDD displays 10 devices. There are five physical drives, and one partition has been assigned on each drive for this configuration. Each SDD device reflects a partition that has been created for a physical drive. Partition 0 stores information about the drive. The operating system masks this partition from the user, but it still exists.

   Note: In a standalone environment, the policy field is optimized. In a cluster environment, the policy field is changed to reserved when a LUN becomes a cluster resource.

```
Total Devices : 10

DEV#:   0  DEVICE NAME: Disk2 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02B12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk2 Part0     OPEN   NORMAL         14          0

DEV#:   1  DEVICE NAME: Disk2 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02B12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk2 Part1     OPEN   NORMAL         94          0

DEV#:   2  DEVICE NAME: Disk3 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02C12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk3 Part0     OPEN   NORMAL         16          0

DEV#:   3  DEVICE NAME: Disk3 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
```

```
 SERIAL: 02C12028

============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk3 Part1     OPEN   NORMAL         94          0

DEV#:   4  DEVICE NAME: Disk4 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02D12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk4 Part0     OPEN   NORMAL         14          0

DEV#:   5  DEVICE NAME: Disk4 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02D22028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk4 Part1     OPEN   NORMAL         94          0

DEV#:   6  DEVICE NAME: Disk5 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02E12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk5 Part0     OPEN   NORMAL         14          0

DEV#:   7  DEVICE NAME: Disk5 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02E12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk5 Part1     OPEN   NORMAL         94          0

DEV#:   8  DEVICE NAME: Disk6 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02F12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk6 Part0     OPEN   NORMAL         14          0

DEV#:   9  DEVICE NAME: Disk6 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02F12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk6 Part1     OPEN   NORMAL         94          0
```

Installing and configuring additional paths

Perform the following steps to install and configure additional paths:

1. Install any additional hardware on the Windows NT server.
2. Install any additional hardware on the supported storage device.
3. Configure the new paths to the server.
4. Restart the Windows NT server. Restarting will ensure correct multipath access to both existing and new storage and to your Windows NT server.
5. Verify that the path is added correctly. See “Verifying additional paths are installed correctly.”

Verifying additional paths are installed correctly

After installing additional paths to SDD devices, verify the following conditions:

- All additional paths have been installed correctly.
- The number of adapters and the number of paths to each storage volume match the updated configuration.
- The Windows disk numbers of all primary paths are labeled as path #0.

Perform the following steps to verify that the additional paths have been installed correctly:

1. Click **Start → Program → Subsystem Device Driver → Subsystem Device Driver Management**. An MS-DOS window opens.
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 path is installed to the previous configuration:

```
Active Adapters :2

Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active
    0  Scsi Port6 Bus0  NORMAL   ACTIVE        188          0     10      10
    1  Scsi Port7 Bus0  NORMAL   ACTIVE        204          0     10      10
```

3. Type **datapath query device** and press **Enter**. The output includes information about any additional devices that were installed. In the example shown in the following output, the output includes information about the new host bus adapter that was assigned:

```
Total Devices : 10

DEV#:   0  DEVICE NAME: Disk2 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02B12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk2 Part0     OPEN   NORMAL          5          0
    1    Scsi Port7 Bus0/Disk7 Part0     OPEN   NORMAL          9          0

DEV#:   1  DEVICE NAME: Disk2 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02B12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk2 Part1     OPEN   NORMAL         32          0
    1    Scsi Port7 Bus0/Disk7 Part1     OPEN   NORMAL         32          0

DEV#:   2  DEVICE NAME: Disk3 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02C12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk3 Part0     OPEN   NORMAL          7          0
    1    Scsi Port7 Bus0/Disk8 Part0     OPEN   NORMAL          9          0

DEV#:   3  DEVICE NAME: Disk3 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02C22028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk3 Part1     OPEN   NORMAL         28          0
    1    Scsi Port7 Bus0/Disk8 Part1     OPEN   NORMAL         36          0

DEV#:   4  DEVICE NAME: Disk4 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02D12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk4 Part0     OPEN   NORMAL          8          0
    1    Scsi Port7 Bus0/Disk9 Part0     OPEN   NORMAL          6          0

DEV#:   5  DEVICE NAME: Disk4 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02D22028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk4 Part1     OPEN   NORMAL         35          0
    1    Scsi Port7 Bus0/Disk9 Part1     OPEN   NORMAL         29          0

DEV#:   6  DEVICE NAME: Disk5 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
```

```
 SERIAL: 02E12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0   Scsi Port6 Bus0/Disk5 Part0      OPEN   NORMAL          6          0
    1   Scsi Port7 Bus0/Disk10 Part0     OPEN   NORMAL          8          0

DEV#:   7  DEVICE NAME: Disk5 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02E22028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk5 Part1     OPEN   NORMAL         24          0
    1   Scsi Port7 Bus0/Disk10 Part1     OPEN   NORMAL         40          0

DEV#:   8  DEVICE NAME: Disk6 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02F12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0   Scsi Port6 Bus0/Disk6 Part0      OPEN   NORMAL          8          0
    1   Scsi Port7 Bus0/Disk11 Part0     OPEN   NORMAL          6          0

DEV#:   9  DEVICE NAME: Disk6 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02F22028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk6 Part1     OPEN   NORMAL         35          0
    1   Scsi Port7 Bus0/Disk11 Part1     OPEN   NORMAL         29          0
```

The definitive way to identify unique volumes on the storage subsystem is by the serial number displayed. The volume appears at the SCSI level as multiple disks (more properly, Adapter/Bus/ID/LUN), but it is the same volume on the ESS. The previous example shows two paths to each partition (path 0: Scsi Port6 Bus0/Disk2, and path 1: Scsi Port7 Bus0/Disk7).

The example shows partition 0 (`Part0`) for each of the devices. This partition stores information about the Windows partition on the drive. The operating system masks this partition from the user, but it still exists. In general, you will see one more partition from the output of the **datapath query device** command than what is being displayed from the Disk Administrator application.

Preferred node path-selection algorithm for the SAN Volume Controller for Cisco MDS 9000

The SAN Volume Controller for Cisco MDS 9000 is a two-controller disk subsystem. SDD distinguishes the paths to a SAN Volume Controller for Cisco MDS 9000 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, then 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 routed only 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.

The following output of the **datapath query device** command shows that the preferred paths are being selected and shows the format of the SAN Volume Controller for Cisco MDS 9000 serial number.

```
Total Devices : 4
DEV#:   0  DEVICE NAME: Disk3 Part0  TYPE: 2145  POLICY: OPTIMIZED
SERIAL: 600567680180000080000000000000001
============================================================
Path#              Adapter/Hard Disk  State  Mode   Select  Errors
    0   Scsi Port3 Bus0/Disk3 Part0    OPEN  NORMAL      2       0
    1   Scsi Port3 Bus0/Disk13 Part0   OPEN  NORMAL      0       0
    2   Scsi Port4 Bus0/Disk23 Part0   OPEN  NORMAL      2       0
    3   Scsi Port4 Bus0/Disk33 Part0   OPEN  NORMAL      0       0
DEV#:   1  DEVICE NAME: Disk3 Part1  TYPE: 2145  POLICY: OPTIMIZED
SERIAL: 600567680180000080000000000000001
 ============================================================
Path#              Adapter/Hard Disk  State  Mode   Select Errors
   0    Scsi Port3 Bus0/Disk3 Part1   OPEN  NORMAL 9679726    20
   1   Scsi Port3 Bus0/Disk13 Part1   OPEN  NORMAL       0     0
   2   Scsi Port4 Bus0/Disk23 Part1   OPEN  NORMAL 1308460     4
   3   Scsi Port4 Bus0/Disk33 Part1   OPEN  NORMAL       0     0
DEV#:   2  DEVICE NAME: Disk11 Part0  TYPE: 2145 POLICY: OPTIMIZED
SERIAL: 600567680180000080000000000000009
 ============================================================
Path#              Adapter/Hard Disk  State   Mode  Select Errors
    0   Scsi Port3 Bus0/Disk11 Part0   OPEN   NORMAL     2      0
    1   Scsi Port3 Bus0/Disk21 Part0   OPEN   NORMAL     0      0
    2   Scsi Port4 Bus0/Disk31 Part0   OPEN   NORMAL     2      0
    3   Scsi Port4 Bus0/Disk41 Part0   OPEN   NORMAL     0      0
 DEV#:   3  DEVICE NAME: Disk11 Part1  TYPE: 2145 POLICY: OPTIMIZED
SERIAL: 600567680180000080000000000000009
============================================================
Path#              Adapter/Hard Disk State  Mode  Select  Errors
    0   Scsi Port3 Bus0/Disk11 Part1 OPEN NORMAL  9965596     15
    1   Scsi Port3 Bus0/Disk21 Part1 OPEN NORMAL        0      0
    2   Scsi Port4 Bus0/Disk31 Part1 OPEN NORMAL 13431178      4
    3   Scsi Port4 Bus0/Disk41 Part1 OPEN NORMAL        0      0
```

Upgrading the SDD

If you attempt to install *over* an existing version of SDD, the installation fails. You must uninstall any previous version of the SDD before installing a new version of SDD.

Perform the following steps to upgrade to a newer SDD version:

1. Uninstall the previous version of SDD. (See “Removing SDD” on page 238 for instructions.)

   Attention: After uninstalling the previous version, you must *immediately* install the new version of SDD to avoid any potential data loss. If you perform a system restart before installing the new version, you may lose access to your assigned volumes.
2. Install the new version of SDD. (See “Installing SDD” on page 230 for instructions.)

Adding or modifying a multipath storage configuration to the supported storage device

This section contains the procedures for adding new storage to an existing configuration in multipath environments.

Reviewing the existing SDD configuration information

Before adding any additional hardware, review the configuration information for the adapters and devices currently on your Windows NT server.

Verify that the number of adapters and the number of paths to each supported storage device volume match the known configuration. Perform the following steps to display information about the adapters and devices:

1. Click **Start → Program → Subsystem Device Driver → Subsystem Device Driver Management**. An MS-DOS window opens.
2. Enter **datapath query adapter** and press **Enter**. The output includes information about all the installed adapters. In the example shown in the following output, two host bus adapters are installed on the Windows NT host server:

```
Active Adapters :2

Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active
    0  Scsi Port6 Bus0  NORMAL   ACTIVE        188          0     10      10
    1  Scsi Port7 Bus0  NORMAL   ACTIVE        204          0     10      10
```

3. Enter **datapath query device** and press **Enter**. In the following example output from an ESS device, four devices are attached to the SCSI path:

```
Total Devices : 2

 DEV#:   0  DEVICE NAME: Disk2 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02B12028
 ============================================================================
 Path#          Adapter/Hard Disk      State     Mode         Select     Errors
    0   Scsi Port5 Bus0/Disk2 Part0     OPEN   NORMAL          4          0
     1   Scsi Port5 Bus0/Disk8 Part0     OPEN   NORMAL          7          0
     2   Scsi Port6 Bus0/Disk14 Part0    OPEN   NORMAL          6          0
     3   Scsi Port6 Bus0/Disk20 Part0    OPEN   NORMAL          5          0

 DEV#:   1  DEVICE NAME: Disk2 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02B12028
 ============================================================================
 Path#          Adapter/Hard Disk       State     Mode     Select     Errors
     0   Scsi Port5 Bus0/Disk2 Part1      OPEN   NORMAL   14792670          0
     1   Scsi Port5 Bus0/Disk8 Part1      OPEN   NORMAL   14799942          0
     2   Scsi Port6 Bus0/Disk14 Part1     OPEN   NORMAL   14926972          0
     3   Scsi Port6 Bus0/Disk20 Part1     OPEN   NORMAL   14931115          0
```

Adding new storage to an existing configuration

Perform the following steps to install additional storage:

1. Install any additional hardware to the supported storage device.
2. Configure the new storage to the server.
3. Restart the Windows NT server. Restarting will ensure correct multipath access to both existing and new storage and to your Windows NT server.
4. Verify that the new storage is added correctly. See “Verifying new storage is installed correctly” on page 237.

Verifying new storage is installed correctly

After adding new storage to an existing configuration, you should verify the following conditions:

- The new storage is correctly installed and configured.
- The number of adapters and the number of paths to each ESS volume match the updated configuration.
- The Windows disk numbers of all primary paths are labeled as path #0.

Perform the following steps to verify that the additional storage has been installed correctly:

1. Click **Start → Program → Subsystem Device Driver → Subsystem Device Driver Management**. An MS-DOS window opens.
2. Enter **datapath query adapter** and press **Enter**. The output includes information about all the installed adapters. In the example shown in the following output, two SCSI adapters are installed on the Windows NT host server:

```
Active Adapters :2

Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active
    0  Scsi Port6 Bus0  NORMAL   ACTIVE        295          0     16      16
    1  Scsi Port7 Bus0  NORMAL   ACTIVE        329          0     16      16
```

3. Enter **datapath query device** and press **Enter**. The output includes information about any additional devices that were installed. In the following example output from an ESS device, the output includes information about the new devices that were assigned:

```
Total Devices : 2

 DEV#:   0  DEVICE NAME: Disk2 Part0  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02B12028
 ============================================================================
 Path#          Adapter/Hard Disk       State     Mode     Select     Errors
     0   Scsi Port5 Bus0/Disk2 Part0      OPEN   NORMAL          4          0
     1   Scsi Port5 Bus0/Disk8 Part0      OPEN   NORMAL          7          0
     2   Scsi Port6 Bus0/Disk14 Part0     OPEN   NORMAL          6          0
     3   Scsi Port6 Bus0/Disk20 Part0     OPEN   NORMAL          5          0

 DEV#:   1  DEVICE NAME: Disk2 Part1  TYPE: 2105E20    POLICY: OPTIMIZED
 SERIAL: 02B12028
 ============================================================================
 Path#          Adapter/Hard   Disk     State     Mode     Select     Errors
     0   Scsi Port5 Bus0/Disk2  Part1     OPEN   NORMAL   14792670          0
     1   Scsi Port5 Bus0/Disk8  Part1     OPEN   NORMAL   14799942          0
     2   Scsi Port6 Bus0/Disk14 Part1     OPEN   NORMAL   14926972          0
     3   Scsi Port6 Bus0/Disk20 Part1     OPEN   NORMAL   14931115          0
```

The definitive way to identify unique volumes on the ESS device is by the serial number displayed. The volume appears at the SCSI level as multiple disks (more properly, Adapter/Bus/ID/LUN), but it is the same volume on the ESS. The previous example shows two paths to each partition (path 0: Scsi Port6 Bus0/Disk2, and path 1: Scsi Port7 Bus0/Disk10).

The example shows partition 0 (`Part0`) for each device. This partition stores information about the Windows partition on the drive. The operating system masks this partition from the user, but it still exists. In general, you will see one more partition from the output of the **datapath query device** command than what is being displayed in the Disk Administrator application.

Removing SDD

Perform the following steps to uninstall SDD on a Windows NT host system:

1. Log on as the administrator user.
2. Click **Start → Settings → Control Panel**. The Control Panel window opens.
3. Double-click **Add/Remove Programs**. The Add/Remove Programs window opens.
4. In the Add/Remove Programs window, select Subsystem Device Driver from the Currently installed programs selection list.
5. Click **Add/Remove**.

Attention:

- After uninstalling the previous version, you must immediately install the new version of SDD to avoid any potential data loss. (See “Installing SDD” on page 230 for instructions.)
- If you perform a system restart and accidentally overwrite the disk signature, you may permanently lose access to your assigned volume. If you do not plan to install the new version of SDD immediately, you need to remove the multipath access to your shared volume. For additional information, refer to the *Multiple-Path Software May Cause Disk Signature to Change* Microsoft article (Knowledge Base Article Number Q293778). This article can be found at the following Web site:

  http://support.microsoft.com

SDD server daemon

The SDD server (also referred to as *sddsrv*) is an integrated component of SDD 1.3.4.x (or later). This component consists of a Windows application daemon that is installed in addition to the SDD device driver. See Chapter 11, “Using the SDD server and the SDDPCM server,” on page 279 for more information about sddsrv.

Verifying that the SDD server has started

After you have installed SDD, verify that the SDD server (sddsrv) has automatically started:

1. Click **Start → Settings → Control Panel**.
2. Double-click **Services**.
3. Look for SDD_Service.

The status of SDD_Service should be *Started* if the SDD server has automatically started.

Starting the SDD server manually

If the SDD server did not start automatically after you performed the SDD installation, you can start sddsrv:

1. Click **Start → Settings → Control Panel**.
2. Double-click **Services**.
3. Select **SDD_Service**.
4. Click **Start**.

Changing to a different port number for the SDD server

To change to a different port number for the SDD server, see “Changing the sddsrv or pcmsrv TCP/IP port number” on page 282.

Stopping the SDD server

You can stop the SDD server by:

- Click **Start → Settings → Control Panel**.
- Double-click **Services**.
- Select **SDD_Service**.
- Click **Stop**.

Displaying the current version of SDD

You can display the current SDD version on a Windows NT host system by viewing the sddpath.sys file properties. Perform the following steps to view the properties of the sddpath.sys file:

1. Click **Start → Programs → Accessories → Windows Explorer**. Windows will open Windows Explorer.
2. In Windows Explorer, go to the %SystemRoot%\system32\drivers directory, where *%SystemRoot%* is:

   %SystemDrive%\winnt for Windows NT.

   If Windows is installed on the C: drive, *%SystemDrive%* is C:. If Windows is installed on E: drive, *%SystemDrive%* is E:
3. Right-click the sddpath.sys file and then click **Properties**. The **sddpath.sys** properties window opens.
4. In the sddpath.sys properties window, click **Version**. The file version and copyright information about the sddpath.sys file is displayed.

Error recovery and retry policy

There are differences in the way that SDD 1.3.4.x (or later) and SDD 1.3.3.1 (or earlier) handles error recovery for Windows NT host systems:

SDD 1.3.3.1 (or earlier) error recovery policy

With SDD 1.3.3.1 (or earlier), error recovery policy is designed to cover a transient kind of error from the user’s applications. The error recovery policy prevents a path from becoming disabled in the event of transient errors.

Note: This policy can cause longer interruption for recovery. The recovery policy halts the I/O activities on functional paths and vpath devices for some time before failing paths are set to DEAD state.

SDD 1.3.4.x (or later) error recovery policy

With SDD 1.3.4.x (or later), error recovery policy is designed to report errors to applications more quickly. With SDD 1.3.4.x (or later), the applications receive failed I/O requests more quickly. This process prevents unnecessary retries, which can cause the I/O activities on good paths and vpath devices to halt for an unacceptable period of time.

Both the SDD 1.3.4.x (or later) and SDD 1.3.3.1 (or earlier) error recovery policies support the following modes of operation:

single-path mode (for ESS only)

A Windows NT host system has only one path that is configured to an ESS LUN.

SDD, in single-path mode, has the following characteristics:

- When an I/O error occurs, SDD retries the I/O operation up to two times.
- With the SDD 1.3.4.x (or later) error recovery policy, SDD returns the failed I/O to the application and sets the state of this failing path to DEAD. SDD driver relies on the SDD server daemon to detect the recovery of the single path. The SDD server daemon recovers this failing path and changes its state to OPEN. (SDD can change a single and failing path into DEAD state.)
- With the SDD 1.3.3.1 (or earlier) error recovery policy, SDD returns the failed I/O to the application and leaves this path in OPEN state (SDD never puts this single path into DEAD state).
- With SDD 1.3.4.x (or later) the SDD server daemon detects the single CLOSE path that is failing and changes the state of this failing path to CLOSE_DEAD. When the SDD server daemon detects a CLOSE_DEAD path that has been recovered, it will change the state of this path to CLOSE. With a single path configured, the SDD vpath device can not be opened if it is the only path in a CLOSE_DEAD state.

multipath mode

The host system has multiple paths that are configured to a supported storage device.

Both the SDD 1.3.4.x (or later) and SDD 1.3.3.1 (or earlier) error recovery policies in multiple-path mode have the following common characteristics:

- If an I/O error occurs on the last operational path to a device, SDD attempts to reuse (or fail back to) a previously failed path.

The SDD 1.3.4.x (or later) error recovery policy in *multiple-path mode* has the following latest characteristics:

- If an I/O error occurs on a path, SDD 1.3.4.x (or later) does not attempt to use the path until three successful I/O operations occur on an operational path.
- If an I/O error occurs consecutively on a path and the I/O error count reaches three, SDD immediately changes the state of the failing path to DEAD.
- Both SDD driver and the SDD server daemon can put a last path into DEAD state, if this path is no longer functional. The SDD server can automatically change the state of this path to OPEN after it is recovered. Alternatively, you can manually change the state of the path to OPEN after it is recovered by using the **datapath set path online** command. Go to “datapath set device path” on page 304 for more information.
- If the SDD server daemon detects that the last CLOSE path is failing, the daemon will change the state of this path to CLOSE_DEAD. The SDD server can automatically recover the path if it detects that it is functional.
- If an I/O fails on all OPEN paths to a supported LUN, SDD returns the failed I/O to the application, and changes the state of all OPEN paths (for failed I/Os) to DEAD, even if some paths did not reach an I/O error count of three.
- If an OPEN path already failed some I/Os, it will not be selected as a retry path.

The SDD 1.3.3.1 (or earlier) error recovery policy in *multiple-path mode* has the following characteristics:

- If an I/O error occurs on a path, SDD 1.3.3.1 does not attempt to use the path until 2000 successful I/O operations occur on an operational path.
- The last path is reserved in OPEN state.
- If an I/O fails on all OPEN paths to a supported LUN, SDD returns the failed I/O to the application and leaves all the paths in OPEN state.
- A failed I/O is retried on all OPEN paths to an ESS LUN even if the OPEN path already failed I/Os.
- SDD changes the failed path from the DEAD state back to the OPEN state after 50 000 successful I/O operations on an operational path.

Using high-availability clustering on an ESS

The following items are required to support the Windows NT operating system on an ESS in a clustering environment:

- SDD 1.2.1 or higher
- Windows NT 4.0 Enterprise Edition with Service Pack 6A
- Microsoft hotfix Q305638 for the clustering environment

Note: SDD does not support I/O load balancing in a Windows NT clustering environment.

Special considerations in the high-availability clustering environment

There are subtle differences in the way that SDD handles path reclamation in a Windows NT clustering environment compared to a nonclustering environment. When the Windows NT server loses a path in a nonclustering environment, the path condition changes from open to dead and the adapter condition changes from active to degraded. The adapter and path condition will not change until the path is made operational again. When the Windows NT server loses a path in a clustering environment, the path condition changes from open to dead and the adapter condition changes from active to degraded. However, after a period of time, the path condition changes back to open and the adapter condition changes back to normal, even if the path has not been made operational again.

The **datapath set adapter # offline** command operates differently in a clustering environment as compared to a nonclustering environment. In a clustering environment, the **datapath set adapter offline** command does not change the condition of the path if the path is active or being reserved. If you issue the command, the following message is displayed: `to preserve access some paths left online`.

Configuring a Windows NT cluster with SDD installed

The following variables are used in this procedure:

server_1 represents the first server with two host bus adapters (HBAs).

server_2 represents the second server with two HBAs.

hba_a represents the first HBA for *server_1*.

hba_b represents the second HBA for *server_1*.

hba_c represents the first HBA for *server_2*.

hba_d represents the second HBA for *server_2*.

Perform the following steps to configure a Windows NT cluster with SDD:

1. Configure LUNs on the ESS as shared for all HBAs on both *server_1* and *server_2*.
2. Connect *hba_a* to the ESS, and restart *server_1*.
3. Click **Start → Programs → Administrative Tools → Disk Administrator**. The Disk Administrator is displayed. Use the Disk Administrator to verify the number of LUNs that are connected to *server_1*.

   The operating system recognizes each additional path to the same LUN as a device.
4. Disconnect *hba_a* and connect *hba_b* to the ESS. Restart *server_1*.
5. Click **Start → Programs → Administrative Tools → Disk Administrator**. The Disk Administrator is displayed. Use the Disk Administrator to verify the number of LUNs that are connected to *server_1*.

   If the number of LUNs that are connected to *server_1* is correct, proceed to 6.

   If the number of LUNs that are connected to *server_1* is incorrect, perform the following steps:

   a. Verify that the cable for *hba_b* is connected to the ESS.

   b. Verify that your LUN configuration on the ESS is correct.

   c. Repeat steps 2 - 5.
6. Install SDD on *server_1*, and restart *server_1*.

   For installation instructions, go to “Installing SDD” on page 230.
7. Connect *hba_c* to the ESS, and restart *server_2*.
8. Click **Start → Programs → Administrative Tools → Disk Administrator**. The Disk Administrator is displayed. Use the Disk Administrator to verify the number of LUNs that are connected to *server_2*.

   The operating system recognizes each additional path to the same LUN as a device.
9. Disconnect *hba_c* and connect *hba_d* to the ESS. Restart *server_2*.
10. Click **Start → Programs → Administrative Tools → Disk Administrator**. The Disk Administrator is displayed. Use the Disk Administrator to verify that the correct number of LUNs are connected to *server_2*.

    If the number of LUNs that are connected to *server_2* is correct, proceed to 11.

    If the number of LUNs that are connected to *server_2* is incorrect, perform the following steps:

    a. Verify that the cable for *hba_d* is connected to the ESS.

    b. Verify your LUN configuration on the ESS.

    c. Repeat steps 7 - 10.
11. Install SDD on *server_2*, and restart *server_2*.

    For installation instructions, go to “Installing SDD” on page 230.
12. Connect both *hba_c* and *hba_d* on *server_2* to the ESS, and restart *server_2*.
13. Use the **datapath query adapter** and **datapath query device** commands to verify the number of LUNs and paths on *server_2*.

14. Click **Start → Programs → Administrative Tools → Disk Administrator**. The Disk Administrator is displayed. Use the Disk Administrator to verify the number of LUNs as online devices. You also need to verify that all additional paths are shown as offline devices.
15. Format the raw devices with NTFS.

    Make sure to keep track of the assigned drive letters on *server_2*.
16. Connect both *hba_a* and *hba_b* on *server_1* to the ESS, and restart *server_1*.
17. Use the **datapath query adapter** and **datapath query device** commands to verify the correct number of LUNs and paths on *server_1*.

    Verify that the assigned drive letters on *server_1* match the assigned drive letters on *server_2*.
18. Restart *server_2*.
    - Install the Microsoft Cluster Server (MSCS) software on *server_1*. When *server_1* is up, install Service Pack 6A (or higher) to *server_1*, and restart *server_1*. Then install hotfix Q305638 and restart *server_1* again.
    - Install the MSCS software on *server_2*. When *server_2* is up, install Service Pack 6A (or higher) to *server_2*, and restart *server_2*. Then install hotfix Q305638 and restart *server_2* again.
19. Use the **datapath query adapter** and **datapath query device** commands to verify 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 volumes and logical volumes for the host server. The secondary server shows only the physical volumes and the logical volumes that it owns.

Search for information about Windows NT clustering.

Making the MoveGroup Service startup type automatic

The MoveGroup Service is shipped with Windows NT 1.3.4.4 or later to enable access to the cluster resources when a movegroup is performed and the primary path is disabled in a cluster environment.

The default startup type of MoveGroup Service is manual. To activate this change, the startup type needs to be automatic. You can change the startup type to automatic as follows:

1. Click **Start → Settings → Control Panel→ Services → SDD MoveGroup Service**.
2. Change Startup type to Automatic.
3. Click **OK**.

After the startup type of MoveGroup Service is changed to automatic, a movegroup of all cluster resources will be performed when a node of the NT cluster is restarted.

Note: The startup type of the MoveGroup Service should the same for both cluster nodes.

Chapter 9. Using SDD on a Windows 2000 host system

This chapter provides procedures for you to install, configure, remove, and use the SDD on a Windows 2000 host system that is attached to a supported storage device.

For updated and additional information not included in this chapter, see the Readme file on the CD-ROM or visit the SDD Web site at

www-1.ibm.com/servers/storage/support/software/sdd.html

Click **Subsystem Device Driver**.

Verifying the hardware and software requirements

You must have the following hardware and software components in order to install SDD:

Hardware

- One or more supported storage devices
- Host system
- For ESS devices: SCSI adapters and cables
- Fibre-channel adapters and cables

Software

- Windows 2000 operating system with Service Pack 2 or higher

  Note: SAN File System may have different Service Pack requirements. Consult the documentation shown in Table 3 on page xxiii for Windows 2000 requirements.
- For ESS devices: SCSI device drivers
- Fibre-channel device drivers

Unsupported environments

SDD does not support the following environments:

- A host system with both a SCSI channel and a fibre-channel connection to a shared LUN.
- Single-path mode during concurrent download and installation of LIC nor during any ESS-concurrent maintenance that impacts the path attachment, such as an ESS host-bay-adapter replacement.
- Support of HBA Symbios SYM8751D has been withdrawn starting with ESS Model 800 and SDD 1.3.3.3.

ESS requirements

To successfully install SDD, ensure that you configure the ESS devices as IBM 2105*xxx* (where *xxx* is the ESS model number) on your Windows 2000 host system.

SAN Volume Controller requirements

To successfully install SDD, ensure that you configure the SAN Volume Controller devices as fibre-channel devices attached to the IBM 2145 on your Windows 2000 host system.

© Copyright IBM Corp. 1999, 2004

SAN Volume Controller for Cisco MDS 9000 requirements

To successfully install SDD, ensure that you configure the SAN Volume Controller for Cisco MDS 9000 devices as fibre-channel devices attached to the IBM 2062 on your Windows 2000 host system.

Host system requirements

To successfully install SDD, your Windows 2000 host system must be an Intel-based system with Windows 2000 Service Pack 2 (or higher) installed. The host system can be a uniprocessor or a multiprocessor system.

To install all components, you must have at least 1 MB (MB equals approximately 1 000 000 bytes) of disk space available on the drive where Windows 2000 is installed.

ESS SCSI requirements

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

- No more than 32 SCSI adapters are attached.
- 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 fibre-channel paths are configured between the host and the subsystem.

  Note: SDD also supports one SCSI adapter on the host system. With single-path access, concurrent download of licensed internal code is supported with SCSI devices. However, the load-balancing and failover features are not available.
- For information about the SCSI adapters that can attach to your Windows 2000 host system, go to the following Web site:

  www.ibm.com/storage/hardsoft/products/ess/supserver.htm

Fibre-channel requirements

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

- Depending on the fabric and supported storage configuration, the number of fibre-channel adapters attached should be less than or equal to 32 / (n * m), where *n* is the number of supported storage ports and *m* is the number of paths that have access to the supported storage device from the fabric.
- A fiber-optic cable connects each fibre-channel adapter to an supported storage port.
- If you need the SDD I/O load-balancing and failover features, ensure that a minimum of two fibre-channel adapters are installed.

Note: You should have at least two fibre-channel adapters to prevent data loss due to adapter hardware failure.

For information about the fibre-channel adapters that can attach to your Windows 2000 host system, go to the following Web site at:
www.ibm.com/storage/hardsoft/products/ess/supserver.htm

Preparing for SDD 1.5.0.0 (or later) installation

Before installing SDD 1.5.0.0 (or later), you must:

Note: If you currently have SDD 1.3.x.x running, then IBM recommends an upgrade to 1.5.0.0 (or later). To upgrade to SDD 1.5.0.0 (or later), see "Upgrading SDD" on page 250.

1. Ensure that all hardware and software requirements are met before you install the SDD. See "Verifying the hardware and software requirements" on page 245 for more information.
2. Configure the supported storage device to your host system. See "Configuring the supported storage device" for more information.
3. Configure the fibre-channel adapters that are attached to your Windows 2000 host system. See "Configuring fibre-channel adapters" for more information.
4. Configure the SCSI adapters that are attached to your Windows 2000 host system. See "Configuring SCSI adapters for ESS devices" on page 248 for more information.

Configuring the supported storage device

Before you install SDD, configure your supported storage device 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 failover features.

For information about configuring your ESS, refer to the *IBM TotalStorage Enterprise Storage Server: Introduction and Planning Guide*.

For information about configuring your SAN Volume Controller, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Configuration Guide*.

For information about configuring your SAN Volume Controller for Cisco MDS 9000, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller for Cisco MDS 9000 Configuration Guide*.

Note: During heavy usage, the Windows 2000 operating system might slow down while trying to recover from error conditions.

Configuring fibre-channel adapters

You must configure the fibre-channel adapters that are attached to your Windows 2000 host system before you install SDD. Follow the adapter-specific configuration instructions to configure the adapters attached to your Windows 2000 host systems. To get the latest recommendation for host adapter settings for ESS, refer to the Enterprise Storage Server interoperability matrix at the following Web site:

www.ibm.com/storage/disk/ess/supserver.htm

To get the latest recommendation for host adapter settings for the SAN Volume Controller, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Host Systems Attachment Guide* and the following Web site:

www.ibm.com/storage/support/2145/

Note: If you are configuring the Emulex LP8000/LP9002 adapter for multipath functions, select **Allow Multiple Paths to SCSI Targets** in the Emulex Configuration Tool panel.

Configuring SCSI adapters for ESS devices

Attention: Failure to disable the BIOS of attached nonstart devices may cause your system to attempt to restart from an unexpected nonstart device.

Before you install and use SDD, you must configure your SCSI adapters. For SCSI adapters that are attached to start devices, ensure that the BIOS for the adapter is enabled. For all other adapters that are attached to nonstart devices, ensure that the BIOS for the adapter is disabled.

Note: When the adapter shares the SCSI bus with other adapters, the BIOS must be disabled.

Installing SDD 1.5.0.0 (or later)

The following section describes how to install SDD 1.5.0.0 (or later) on your system.

Use the following default settings for local policies/security:

Policy	Setting
unsigned driver installation behavior	Not defined
unsigned non-driver installation behavior	Not defined

These default settings on a Windows 2000 machine are documented in the Microsoft Web site:

www.microsoft.com/technet/treeview/default.asp?url=/technet/security/issues/W2kCCSCG/W2kSCGca.asp

If you do not use the default setting for unsigned non-driver installation behavior, use the Silently Succeed setting.

Note: Ensure that SDD is installed *before* adding additional paths to a device. Otherwise, the Windows 2000 server could lose the ability to access existing data on that device.

Perform the following steps to install SDD 1.5.0.0 (or later) on your system:

1. Log on as the administrator user.
2. Insert the SDD installation CD-ROM into the selected drive.
3. Start the Windows 2000 Explorer program.
4. Double-click the CD-ROM drive. A list of all the installed directories on the compact disc is displayed.
5. Double-click the \win2k\IBMsdd directory (or your installation subdirectory)
6. Run the setup.exe program. The setup program starts.

   Tip:

   - If you have previously installed a 1.3.1.1 (or earlier) version of SDD, you will see an *"Upgrade?"* question while the setup program is running. You should answer **y** to this question to continue the installation. Follow the displayed setup instructions to complete the installation.
   - If you currently have SDD 1.3.1.2 or 1.3.2.x installed on your Windows 2000 host system, answer **y** to the *"Upgrade?"* question.

7. When the setup program is finished, you will be asked if you want to reboot. If you answer **y**, setup will restart your Windows 2000 system immediately. Otherwise setup will exit, and you will need to manually restart your Windows 2000 system to activate the new installation.
8. If this is a new installation:
   a. Shut down your Windows 2000 host system.
   b. Reconnect all cables that connect the host bus adapters and the supported storage devices if needed.
   c. Restart your Windows 2000 host system.
9. If this is an upgrade, restart your Windows 2000 host system.

After completing the installation procedures and when you log on again, your Program menu will include a **Subsystem Device Driver** entry containing the following selections:

1. Subsystem Device Driver management
2. SDD Technical Support Web site
3. README

Notes:

1. You can verify that SDD has been successfully installed by issuing the **datapath query device** command. If the command executes, SDD is installed.

   You can also use the following operation to verify that SDD 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.

      IBM 2105xxx SDD Disk Device : indicates ESS devices connected to Windows 2000 host

      IBM 2105xxx SCSI Disk Device : indicates path access to ESS devices

   Figure 6 on page 250 shows six ESS devices connected to the host and four paths to each of the ESS devices. The Device manager shows six IBM 2105xxx SDD Disk Devices and 24 IBM 2105xxx SCSI Disk Devices.

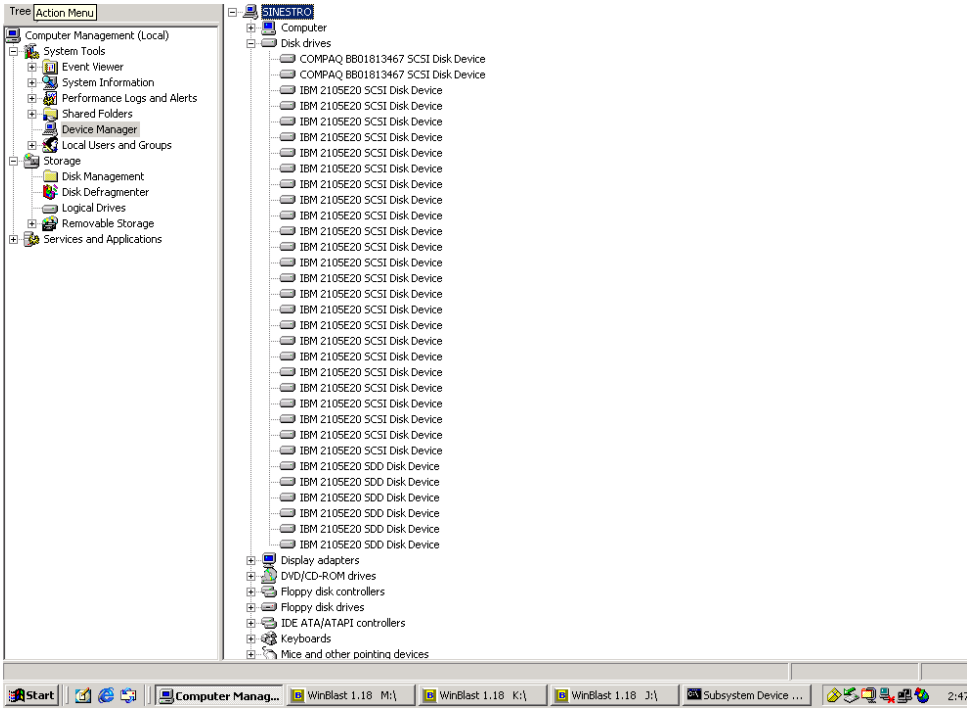

Figure 6. Example showing ESS devices to the host and path access to the ESS devices in a successful SDD installation on a Windows 2000 host system

2. You can also verify the current version of SDD. For more information, go to "Displaying the current version of SDD."

Displaying the current version of SDD

You can display the current version of SDD on a Windows 2000 host system by viewing the sddbus.sys file properties. Perform the following steps to view the properties of sddbus.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%\winnt for Windows 2000.

   If Windows is installed on the C: drive, *%SystemDrive%* is C:. If Windows is installed on E: drive, *%SystemDrive%* is E:
3. Right-click the sddbus.sys file, and then click **Properties**. The sddbus.sys properties window opens.
4. In the sddbus.sys properties window, click **Version**. The file version and copyright information about the sddbus.sys file is displayed.

Upgrading SDD

IBM recommends that you perform the upgrade to SDD 1.4.0.0 (or later) if you currently have SDD 1.3.x.x installed on your Windows 2000 host system.

Follow the instructions given in "Installing SDD 1.5.0.0 (or later)" on page 248 to upgrade SDD.

Note:

- You can verify that SDD has been successfully installed by issuing the **datapath query device** command. If the command executes, SDD is installed.
- You can also verify the current version of SDD. For more information, go to “Displaying the current version of SDD” on page 250.

Upgrading to SDD 1.4.0.0 (or later) in a two-node cluster environment

If you have SDD 1.3.1.2, SDD 1.3.2.x, SDD 1.3.3.x or SDD 1.3.4.x installed on your Windows 2000 host system, perform the following steps for upgrading to SDD 1.4.0.0 (or later) in a two-node cluster environment:

1. Move all cluster resources from node A to node B.
2. Follow the instructions from the “Upgrading SDD” on page 250 on node A.
3. When node A is up, move all resources from node B to node A.
4. Follow the instructions from the “Upgrading SDD” on page 250 on node B.

Remote boot support for ESS

The following procedures describe how to implement remote boot support for ESS devices connected to a fibre-channel host bus adapter.

Note: Support for remote boot from ESS devices connected to a SCSI adapter is not available.

Booting from an ESS device with Windows 2000 and SDD 1.5.0.0 (or later) using QLOGIC HBA.

Perform the following steps to install Subsystem Device Driver :

1. Configure the ESS and SAN environment.
2. Obtain the World Wide name of the Qlogic HBA you are going to boot from. This can be obtained by entering the Qlogic BIOS through **CRTL+Q** in the Adapter Setting Panel.

   Note: The second HBA will not be configured at this time.
3. Boot the server you are setting up boot support for. Ensure that there is only 1 path from your QLogic HBA to ESS storage. Enter **CTRL+Q** to enter the QLOGIC BIOS Fast Utility.
4. Select the boot support HBA.
5. Select **Configuration Settings**.
6. Use the Host Adapter Setting Panel and enable the BIOS for the adapter.
7. Use the Selectable Boot Settings panel and enable Selectable Boot.
8. Select the first (primary) boot and press **Enter**.
9. Select **IBM device** and press **Enter**.
10. At the Select LUN prompt, select the first supported LUN, which is LUN 0.
11. Save changes and reboot system using bootable Windows 2000 diskettes or CD-ROM.
12. At the first Windows 2000 installation screen, press **F6** to install a third party device.
13. Select **S** to specify an additional device.
14. Insert the floppy disk that contains the Qlogic HBA driver and press **Enter**.

15. Continue to install Windows 2000. Select the first ESS volume seen by Qlogic HBA as the device on which to install Windows 2000.
16. Install the Windows 2000 Service Pack.
17. Install SDD and reboot. You may be asked by the system to reboot 1 more time.
18. Shut down the system.
19. Connect fibre-channel cables from the other Qlogic HBA to ESS storage.
20. Ensure that the BIOS of this adapter is disabled.
21. Add multipaths to ESS.
22. Restart the system.

Limitations when booting from an ESS device on a Windows 2000 host

The following limitations apply when booting from an ESS device on a Windows 2000 host:

1. You cannot use the same HBA as both the ESS boot device and a clustering adapter. This is a Microsoft physical limitation.
2. If you reboot a system with adapters while the primary path is in failed state , you must:
   a. Manually disable the BIOS on the first adapter
   b. Manually enable the BIOS on the second adapter.
3. You cannot enable the BIOS for both adapters at the same time. If the BIOS for both adapters is enabled at the same time and there is path failure on the primary adapter, the system will error with INACCESSIBLE_BOOT_DEVICE upon reboot.

Uninstalling SDD

Perform the following steps to uninstall SDD on a Windows 2000 host system:

1. Shut down your Windows 2000 host system.
2. Ensure that there is a single-path connection from the system to the storage device.
3. Turn on your Windows 2000 host system.
4. Log on as the administrator user.
5. Click **Start → Settings → Control Panel**. The Control Panel opens.
6. Double-click **Add/Remove Programs**. The Add/Remove Programs window opens.
7. In the Add/Remove Programs window, select the Subsystem Device Driver from the currently installed programs selection list.
8. Click **Add/Remove**.
9. Restart your system.

Attention:

- You must install SDD 1.5.0.0 (or later) immediately after performing a system restart to avoid any potential data loss. Go to “Installing SDD 1.5.0.0 (or later)” on page 248 for instructions.
- If you are not planning to reinstall the Subsystem Device Driver 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.

Removing SDD in a two-node cluster environment

IBM recommends the following steps if you intend to remove the multipathing functions from an supported storage device in a two-node cluster environment.

Perform the following steps to remove SDD 1.4.0.0 (or later) in a two-node cluster environment:

1. Move all cluster resources from node A to node B.
2. Ensure that there is 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.
   d. Remove multiple supported storage port access, if applicable.
3. Uninstall SDD. See “Uninstalling SDD” on page 252 for details.
4. Restart your system.
5. Move all cluster resources from node B to node A.
6. Perform steps 2 - 5 on node B.

SDD server daemon

The SDD server (also referred to as sddsrv) is an integrated component of SDD 1.3.4.1 (or later). This component consists of a Windows application daemon that is installed in addition to the SDD device driver. See Chapter 11, “Using the SDD server and the SDDPCM server,” on page 279 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:

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**. The status of SDD Service should be *Started* if the SDD server has automatically started.

Starting the SDD server manually

If the SDD server did not start automatically after you performed the SDD installation, you can use the following process 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**.

Changing to a different port number for the SDD server

To change to a different port number for the SDD server, see “Changing the sddsrv or pcmsrv TCP/IP port number” on page 282.

Stopping the SDD server

To stop the SDD server:

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**.

Adding paths to SDD devices

To activate SDD, you need to restart your Windows 2000 system after it is installed.

Attention: Ensure that SDD is installed before you add additional paths to a device. Otherwise, the Windows 2000 server could lose the ability to access existing data on that device.

Before adding any additional hardware, review the configuration information for the adapters and devices currently on your Windows 2000 server. Perform the following steps to display information about the adapters and devices:

1. You must log on as an administrator user to have access to the Windows 2000 Computer Management.
2. Click **Start → Program → Subsystem Device Driver → Subsystem Device Driver Management**. An MS-DOS window opens.
3. Enter **datapath query adapter** and press **Enter**. The output includes information about all the installed adapters. In the example shown in the following output, one host bus adapter is installed:

```
Active Adapters :1

Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active
    0  Scsi Port4 Bus0  NORMAL   ACTIVE        592          0     6      6
```

4. Enter **datapath query device** and press **Enter**. In the following example showing ESS device output, eight devices are attached to the SCSI path:

```
Total Devices : 6


DEV#:   0  DEVICE NAME: Disk1 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06D23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk1 Part0     OPEN   NORMAL        108          0

DEV#:   1  DEVICE NAME: Disk2 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06E23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk2 Part0     OPEN   NORMAL         96          0

DEV#:   2  DEVICE NAME: Disk3 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06F23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk3 Part0     OPEN   NORMAL         96          0

DEV#:   3  DEVICE NAME: Disk4 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07023922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk4 Part0     OPEN   NORMAL         94          0

DEV#:   4  DEVICE NAME: Disk5 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07123922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk5 Part0     OPEN   NORMAL         90          0

DEV#:   5  DEVICE NAME: Disk6 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07223922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk6 Part0     OPEN   NORMAL         98          0
```

Activating additional paths

Perform the following steps to activate additional paths to a vpath device:

1. Install any additional hardware on the Windows 2000 server or the ESS.
2. Click **Start → Program → Administrative Tools → Computer Management**.
3. Click **Device Manager**.
4. Right-click **Disk drives**.
5. Click **Scan for hardware changes**.
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 SDD devices, verify that the additional paths have been installed correctly.

Perform the following steps to verify that the additional paths have been installed correctly:

1. Click **Start → Program → Subsystem Device Driver → Subsystem Device Driver Management**. An MS-DOS window opens.

2. Enter **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 host bus adapter has been installed:

```
Active Adapters :2

Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active
    0  Scsi Port1 Bus0  NORMAL   ACTIVE       1325          0      8       8
    1  Scsi Port2 Bus0  NORMAL   ACTIVE       1312          0      8       8
```

3. Enter **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 host bus adapter and the new device numbers that were assigned. For ESS devices, the following output is displayed:

```
Total Devices : 6


DEV#:   0  DEVICE NAME: Disk1 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06D23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk1 Part0     OPEN   NORMAL        108          0
    1    Scsi Port5 Bus0/Disk1 Part0     OPEN   NORMAL         96          0

DEV#:   1  DEVICE NAME: Disk2 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06E23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk2 Part0     OPEN   NORMAL         96          0
    1    Scsi Port5 Bus0/Disk2 Part0     OPEN   NORMAL         95          0

DEV#:   2  DEVICE NAME: Disk3 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06F23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk3 Part0     OPEN   NORMAL         96          0
    1    Scsi Port5 Bus0/Disk3 Part0     OPEN   NORMAL         94          0

DEV#:   3  DEVICE NAME: Disk4 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07023922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk4 Part0     OPEN   NORMAL         94          0
    1    Scsi Port5 Bus0/Disk4 Part0     OPEN   NORMAL         96          0

DEV#:   4  DEVICE NAME: Disk5 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07123922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk5 Part0     OPEN   NORMAL         90          0
    1    Scsi Port5 Bus0/Disk5 Part0     OPEN   NORMAL         99          0

DEV#:   5  DEVICE NAME: Disk6 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07223922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk6 Part0     OPEN   NORMAL         98          0
    1    Scsi Port5 Bus0/Disk6 Part0     OPEN   NORMAL         79          0
```

Preferred Node path-selection algorithm for the SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000

The SAN Volume Controller and the SAN Volume Controller for Cisco MDS 9000 is a two-controller disk subsystem. SDD distinguishes the paths to a SAN Volume Controller or SAN Volume Controller for Cisco MDS 9000 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, then 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.

The following output of the **datapath query device** command shows that the preferred paths are being selected and shows the format of the SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 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
```

Error recovery and retry policy

With SDD 1.4.0.0 (or later), error recovery policy is designed to report failed I/O requests to applications more quickly. This process prevents unnecessary retries, which can cause the I/O activities on good paths and vpath devices to halt for an unacceptable period of time.

SDD 1.4.0.0 (or later) error recovery policies support the following modes of operation:

single-path mode (for ESS only)
An Windows 2000 host system has only one path that is configured to an ESS logical unit number (LUN).

SDD, in single-path mode, has the following characteristics:

- When an I/O error occurs, SDD retries the I/O operation up to two times.
- With the SDD 1.4.0.0 (or later) error recovery policy, SDD returns the failed I/O to the application and sets the state of this failing path to DEAD. SDD driver relies on the SDD server daemon to detect the recovery of the single path. The SDD server daemon recovers this failing path and changes its state to OPEN. (SDD can change a single and failing path into DEAD state.)
- With SDD 1.4.0.0 (or later), the SDD server daemon detects the single CLOSE path that is failing and changes the state of this failing path to CLOSE_DEAD. When the SDD server daemon detects a CLOSE_DEAD path recovered, it changes the state of this path to CLOSE. With a single path configured, the SDD vpath device can not be opened if it is the only path in a CLOSE_DEAD state.

multipath mode

The host system has multiple paths that are configured to a supported storage device LUN.

SDD 1.4.0.0 (or later) error recovery policies in multiple-path mode have the following common characteristics:

- If an I/O error occurs on the last operational path to a device, SDD attempts to reuse (performs a failback operation to return to) a previously failed path.

The SDD 1.4.0.0 (or later) error recovery policy in multipath mode has the following latest characteristics:

- SDD 1.4.0.0 (or later) does not attempt to use the path until three successful I/O operations occur on an operational path.
- If an I/O error occurs consecutively on a path and the I/O error count reaches three, SDD immediately changes the state of the failing path to DEAD.
- Both SDD driver and the SDD server daemon can put a last path into DEAD state, if this path is no longer functional. The SDD server can automatically change the state of this path to OPEN after it is recovered. Alternatively, you can manually change the state of the path to OPEN after it is recovered by using the **datapath set path online** command. Go to “datapath set device path” on page 304 for more information.
- If an I/O fails on all OPEN paths to an storage device LUN, SDD returns the failed I/O to the application and changes the state of all OPEN paths (for failed I/Os) to DEAD, even if some paths did not reach an I/O error count of three.
- If an OPEN path already failed some I/Os, it will not be selected as a retry path.

Note: When a path failover does not work with the QLogic card, you need to verify that the Target Enabled bit is set in the QLogic BIOS.

Press **Ctrl+Q** at boot time to change the QLogic BIOS settings.

Refer to the following Web site for recommended QLogic BIOS settings:

http://publibfp.boulder.ibm.com/epubs/pdf/f2bhs00.pdf

Support for Windows 2000 clustering

SDD 1.3.0.0 or higher is required to support Windows 2000 clustering. SDD 1.3.0.0 or higher does not support I/O load balancing in a Windows 2000 clustering environment.

When running Windows 2000 clustering, clustering failover might not occur when the last path is being removed from the shared resources. See *Microsoft article Q294173* for additional information. Windows 2000 does not support dynamic disks in the MSCS environment.

Special considerations in the Windows 2000 clustering environment

There are subtle differences in the way that SDD handles path reclamation in a Windows 2000 clustering environment compared to a nonclustering environment. When the Windows 2000 server loses a path in a nonclustering environment, the path condition changes from open to dead and the adapter condition changes from active to degraded. The adapter and path condition will not change until the path is made operational again. When the Windows 2000 server loses a path in a clustering environment, the path condition changes from open to dead and the adapter condition changes from active to degraded. However, after a period of time, the path condition changes back to open and the adapter condition changes back to normal, even if the path has not been made operational again.

Note: The adapter goes to DEGRAD state when there are active paths left on the adapter. It goes to FAILED state when there are no active paths.

The **datapath set adapter # offline** command operates differently in a clustering environment as compared to a nonclustering environment. In a clustering environment, the **datapath set adapter offline** command does not change the condition of the path if the path is active or being reserved. If you issue the command, the following message is displayed: `to preserve access some paths left online`.

Configuring a Windows 2000 cluster with SDD installed

The following variables are used in this procedure:

server_1	Represents the first server with two host bus adapters (HBAs).
server_2	Represents the second server with two HBAs.
hba_a	Represents the first HBA for *server_1*.
hba_b	Represents the second HBA for *server_1*.
hba_c	Represents the first HBA for *server_2*.
hba_d	Represents the second HBA for *server_2*.

Perform the following steps to configure a Windows 2000 cluster with SDD:

1. Configure LUNs on the storage device as shared for all HBAs on both *server_1* and *server_2*.
2. Connect *hba_a* to the storage device, and restart *server_1*.
3. Click **Start → Programs → Administrative Tools → Computer Management**. The Computer Management window opens. From the Computer Management window, select Storage and then Disk Management to work with the storage devices attached to the host system.

The operating system will recognize each additional path to the same LUN as a device.

4. Disconnect *hba_a* and connect *hba_b* to the ESS. Restart *server_1*.
5. Click **Start → Programs → Administrative Tools → Computer Management**. The Computer Management window opens. From the Computer Management window, select Storage and then Disk Management to verify the correct number of LUNs that are connected to *server_1*.

   If the number of LUNs that are connected to *server_1* is correct, proceed to 6 on page 261.

   If the number of LUNs that are connected to *server_1* is incorrect, perform the following steps:

   a. Verify that the cable for *hba_b* is connected to the ESS.
   b. Verify your LUN configuration on the storage device.
   c. Repeat steps 2 - 5.

6. Install SDD on *server_1*, and restart *server_1*.

   For installation instructions, go to “Installing SDD 1.5.0.0 (or later)” on page 248 section.
7. Connect *hba_c* to the ESS, and restart *server_2*.
8. Click **Start → Programs → Administrative Tools → Computer Management**. The Computer Management window opens. From the Computer Management window, select Storage and then Disk Management to verify the correct number of LUNs that are connected to *server_2*.

   The operating system will see each additional path to the same LUN as a device.
9. Disconnect *hba_c* and connect *hba_d* to the ESS. Restart *server_2*.
10. Click **Start → Programs → Administrative Tools → Computer Management**. The Computer Management window is displayed. From the Computer Management window, select Storage and then Disk Management to verify the correct number of LUNs that are connected to *server_2*.

    If the number of LUNs that are connected to *server_2* is correct, proceed to 11.

    If the number of LUNs that are connected to *server_2* is incorrect, perform the following steps:

    a. Verify that the cable for *hba_d* is connected to the ESS.

    b. Verify your LUN configuration on the storage device.

    c. Repeat steps 7 - 10.
11. Install SDD on *server_2*, and restart *server_2*.

    For installation instructions, go to “Installing SDD 1.5.0.0 (or later)” on page 248.
12. Connect both *hba_c* and *hba_d* on *server_2* to the ESS, and restart *server_2*.
13. Use the **datapath query adapter** and **datapath query device** commands to verify the correct number of LUNs and paths on *server_2*.
14. Click **Start → Programs → Administrative Tools → Computer Management**. The Computer Management window opens. From the Computer Management window, select Storage and then Disk Management to verify that the actual number of LUNs as online devices is correct.
15. Format the raw devices with NTFS.

    Make sure to keep track of the assigned drive letters on *server_2*.
16. Connect both *hba_a* and *hba_b* on *server_1* to the ESS, and restart *server_1*.
17. Use the **datapath query adapter** and **datapath query device** commands to verify the correct number of LUNs and paths on *server_1*.

    Verify that the assigned drive letters on *server_1* match the assigned drive letters on *server_2*.

18. Restart *server_2*.
    - Install the MSCS software on *server_1*, restart *server_1*, reapply Service Pack 2 or higher to *server_1*, and restart *server_1* again.
    - Install the MSCS software on *server_2*, restart *server_2*, reapply Service Pack 2 to *server_2*, and restart *server_2* again.
19. Use the **datapath query adapter** and **datapath query device** commands to verify 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.

Information about installing a Windows 2000 cluster can be found at:
www.microsoft.com/windows2000/techinfo/planning/server/clustersteps.asp

Chapter 10. Using SDD on a Windows Server 2003 host system

This chapter provides procedures for you to install, configure, remove, and use the SDD on a Windows Server 2003 host system that is attached to a supported storage device.

SDD supports both 32-bit and 64-bit environments running Windows Server 2003. For the Windows 2003 Server 32-bit environment, install the package from the \win2k3\i386\IBMsdd directory of the SDD CD-ROM. For the Windows 2003 Server 64-bit environment, install the package from the \win2k3\IA64\IBMsdd directory of the SDD CD-ROM.

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

www-1.ibm.com/servers/storage/support/software/sdd.html

Click **Subsystem Device Driver**.

Verifying the hardware and software requirements

You must have the following hardware and software components in order to install SDD:

Hardware

- Supported storage devices
- Host system
- SCSI adapters and cables
- Fibre-channel adapters and cables

Software

- Windows Server 2003 operating system Standard or Enterprise edition.
- Device driver for SCSI or fibre-channel adapters

Unsupported environments

SDD does not support the following environments:

- A host system with both a SCSI channel and a fibre-channel connection to a shared LUN.
- Windows Server 2003 operating system or a paging file on an SDD-controlled multipath device (that is, SDD does not support boot from the ESS device).
- Single-path mode during concurrent download and installation of LIC nor during any ESS concurrent maintenance that impacts the path attachment, such as an ESS host-bay-adapter replacement.
- SDD is not supported on the Windows Server 2003 Web edition.

ESS requirements

To successfully install SDD, make sure that you configure the ESS devices as IBM 2105*xxx* (where *xxx* is the ESS model number) on your Windows Server 2003 host system.

© Copyright IBM Corp. 1999, 2004

Host system requirements

To successfully install SDD, your Windows Server 2003 host system must be an Intel-based system with Windows Server 2003 installed. The host system can be a uniprocessor or a multiprocessor system.

To install all components, you must have at least 1 MB (MB equals approximately 1 000 000 bytes) of disk space available on the drive where Windows Server 2003 is installed.

SCSI requirements

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

- No more than 32 SCSI adapters are attached.
- 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.

  Note: SDD also supports one SCSI adapter on the host system. With single-path access, concurrent download of licensed internal code is supported with SCSI devices. However, the load-balancing and failover features are not available.
- For information about the SCSI adapters that can attach to your Windows Server 2003 host system, go to the following Web site:

  www.ibm.com/storage/hardsoft/products/ess/supserver.htm

Fibre-channel requirements

To use the SDD 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 an ESS port.
- If you need the SDD 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 ESS 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 2003 host system, go to the following Web site at:
www.ibm.com/storage/hardsoft/products/ess/supserver.htm

Preparing for SDD 1.5.0.0 (or later) installation

Note: If you have Windows 2000 server running and SDD 1.3.x.x already installed and you want to upgrade to Windows Server 2003, you should:

1. Upgrade SDD to 1.5.0.0 (or later).
2. Upgrade Windows 2000 server to Windows Server 2003.

Before installing SDD 1.5.0.0 (or later), you must:

1. Ensure that all hardware and software requirements are met before you install SDD. See “Verifying the hardware and software requirements” on page 263 for more information.
2. Configure the ESS to your host system. See “Configuring the ESS” for more information.
3. Configure the fibre-channel adapters that are attached to your Windows Server 2003 host system. See “Configuring fibre-channel adapters” for more information.
4. Configure the SCSI adapters that are attached to your Windows Server 2003 host system. See “Configuring SCSI adapters” for more information.
5. Uninstall any previously installed version of SDD on your host system.

   For SDD uninstallation and installation instructions, go to:

   a. “Uninstalling SDD” on page 269
   b. “Installing SDD 1.5.0.0 (or later)” on page 266

Configuring the ESS

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 failover features.

For information about configuring your ESS, refer to the *IBM TotalStorage Enterprise Storage Server: Introduction and Planning Guide*.

Note: During heavy usage, the Windows Server 2003 operating system might slow down while trying to recover from error conditions.

Configuring the SAN Volume Controller

Before you install SDD, configure your supported storage device 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 failover features.

For information about configuring your SAN Volume Controller, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Configuration Guide*.

Note: During heavy usage, the Windows Server 2003 operating system might slow down while trying to recover from error conditions.

Configuring fibre-channel adapters

You must configure the fibre-channel adapters that are attached to your Windows Server 2003 host system before you install SDD. Follow the adapter-specific configuration instructions to configure the adapters attached to your Windows Server 2003 host systems.

SDD supports the Emulex HBA with full-port driver. When you configure the Emulex HBA for multipath functions, select **Allow Multiple Paths to SCSI Targets** in the Emulex Configuration Tool panel.

Configuring SCSI adapters

Attention: Failure to disable the BIOS of attached nonstart devices may cause your system to attempt to restart from an unexpected nonstart device.

Before you install and use SDD, you must configure your SCSI adapters. For SCSI adapters that are attached to start devices, ensure that the BIOS for the adapter is enabled. For all other adapters that are attached to nonstart devices, ensure that the BIOS for the adapter is disabled.

Note: When the adapter shares the SCSI bus with other adapters, the BIOS must be disabled.

Installing SDD 1.5.0.0 (or later)

The following section describes how to install SDD 1.5.0.0 (or later) on your system.

Use the following default settings for local policies/security:

Policy	Setting
unsigned driver installation behavior	Not defined
unsigned non-driver installation behavior	Not defined

These default settings on a Windows 2000 machine are documented in the Microsoft Web site:

www.microsoft.com/technet/treeview/default.asp?url=/technet/security/issues/W2kCCSCG/W2kSCGca.asp

If you do not use the default setting for unsigned non-driver installation behavior, use the Silently Succeed setting.

Note:

- Ensure that SDD is installed *before* adding additional paths to a device. Otherwise, the Windows Server 2003 server could lose the ability to access existing data on that device.

Perform the following steps to install SDD 1.5.0.0 (or later) 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 SDD is installed.
2. Log on as the administrator user.
3. Insert the SDD installation CD-ROM into the selected drive.
4. Start the Windows Server 2003 Explorer program.
5. Select the CD-ROM drive. A list of all the installed directories on the compact disc is displayed.
6. Select the \win2k\IBMsdd directory (or your installation subdirectory).
7. Run the setup.exe program. The setup program starts.
8. Shut down your Windows Server 2003 host system.
9. Connect additional cables to your storage if needed.
10. Restart your Windows Server 2003 host system.

After completing the installation procedures and when you log on again, you will see a **Subsystem Device Driver** entry in your Program menu containing the following selections:

1. Subsystem Device Driver Management
2. SDD Technical Support Web site

3. README

Notes:

1. You can verify that SDD has been successfully installed by issuing the **datapath query device** command. If the command executes, SDD is installed.

   You can also use the following operation to verify that SDD 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.

      IBM 2105xxx SDD Disk Device : indicates ESS devices connected to Windows Server 2003 host

      IBM 2105xxx SCSI Disk Device : indicates path access to ESS devices

   In Figure 7, there are six ESS devices connected to the host and four paths to each of the ESS devices. The Device manager shows six IBM 2105xxx SDD Disk Devices and 24 IBM 2105xxx SCSI Disk Devices.

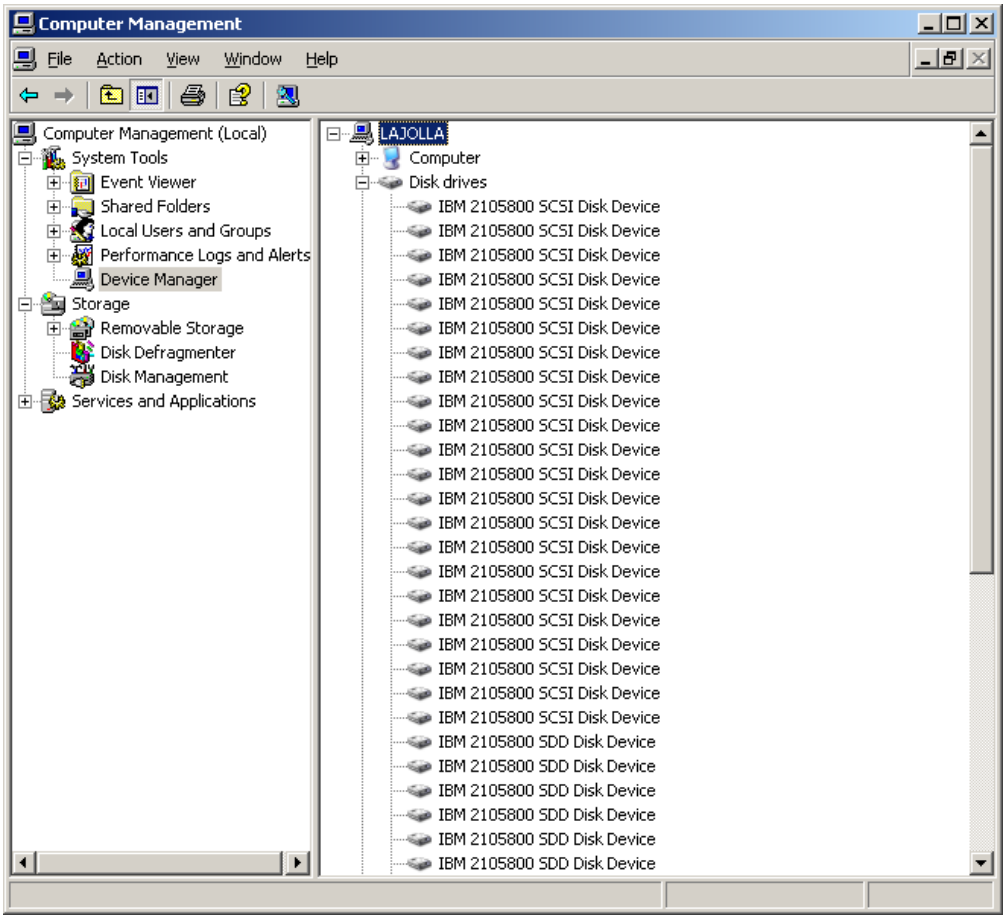

Figure 7. Example showing ESS devices to the host and path access to the ESS devices in a successful SDD installation on a Windows Server 2003 host system

2. You can also verify the current version of SDD. For more information, go to "Displaying the current version of SDD" on page 268.

Displaying the current version of SDD

You can display the current version of SDD on a Windows Server 2003 host system by viewing the sddbus.sys file properties. Perform the following steps to view the properties of sddbus.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 2003.

   If Windows is installed on the C: drive, *%SystemDrive%* is C:. If Windows is installed on E: drive, *%SystemDrive%* is E:
3. Right-click the sddbus.sys file, and then click **Properties**. The sddbus.sys properties window opens.
4. In the sddbus.sys properties window, click **Version**. The file version and copyright information about the sddbus.sys file is displayed.

Upgrading SDD

Use the following procedure to upgrade SDD.

Upgrading from a Windows NT host system to Windows Server 2003

Use the following procedure to upgrade SDD to a Windows Server 2003 host:

1. Uninstall SDD from the Windows NT host system. See “Removing SDD” on page 238.
2. Shut down the system.
3. Disconnect all cables that allow the Windows NT host to access to the supported storage devices.
4. Restart the system.
5. Perform the Windows NT to Windows Server 2003 upgrade according to your migration plans.
6. After your host upgrade is complete, install Windows Server 2003-supported HBA drivers.
7. Enable a single-path access from your server to the supported storage device.
8. Restart your host.
9. Install the latest version of SDD for Windows 2003. See “Installing SDD 1.5.0.0 (or later)” on page 266.
10. Reboot the system, enabling additional paths to the supported storage device.

Uninstalling SDD

Attention:

- You must install SDD 1.5.0.0 (or later) immediately after performing a system restart to avoid any potential data loss. Go to “Installing SDD 1.5.0.0 (or later)” on page 266 for instructions.
- If you are not planning to reinstall the Subsystem Device Driver 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.

Perform the following steps to uninstall SDD on a Windows Server 2003 host system:

1. Log on as the administrator user.
2. Click **Start → Settings → Control Panel**. The Control Panel opens.
3. Double-click **Add/Remove Programs**. The Add/Remove Programs window opens.
4. In the Add/Remove Programs window, select **Subsystem Device Driver** from the currently installed programs selection list.
5. Click **Add/Remove**.
6. Shut down your Windows Server 2003 host system after the uninstallation has completed.
7. Change the zoning configuration or cable connections to ensure that there is only single-path connection from the system to the storage device.
8. Power on your Windows Server 2003 host system.

Removing SDD in a two-node cluster environment

IBM recommends the following steps if you intend to remove the multipathing functions to an ESS device in a two-node cluster environment.

Perform the following steps to remove SDD 1.5.0.0 (or later) 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 ESS port access, if applicable.
3. Uninstall SDD. See “Uninstalling SDD” for instructions.
4. Restart your system.
5. Move all cluster resources from node B to node A.
6. Perform steps 2 - 5 on node B.

SDD server daemon

The SDD server (also referred to as sddsrv) is an integrated component of SDD 1.5.0.0 (or later). This component consists of a Windows application daemon that is installed in addition to the SDD device driver. See Chapter 11, “Using the SDD server and the SDDPCM server,” on page 279 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:

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**. The status of SDD Service should be *Started* if the SDD server has automatically started.

Starting the SDD server manually

If the SDD server did not start automatically after you performed the SDD installation, you can 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**.

Changing to a different port number for the SDD server

To change to a different port number for the SDD server, see “Changing the sddsrv or pcmsrv TCP/IP port number” on page 282.

Stopping the SDD server

To stop the SDD server:

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**.

Adding paths to SDD devices

To activate SDD, you need to restart your Windows Server 2003 system after it is installed.

Attention: Ensure that SDD is installed before you add additional paths to a device. Otherwise, the Windows Server 2003 server could lose the ability to access existing data on that device.

Before adding any additional hardware, review the configuration information for the adapters and devices currently on your Windows Server 2003 server. Perform the following steps to display information about the adapters and devices:

1. You must log on as an administrator user to have access to the Windows Server 2003 Computer Management.
2. Click **Start → Program → Subsystem Device Driver → Subsystem Device Driver Management**. An MS-DOS window opens.
3. Enter **datapath query adapter** and press **Enter**. The output includes information about all the installed adapters. In the example shown in the following output, one host bus adapter is installed:

```
Active Adapters :1

Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active
    0  Scsi Port4 Bus0  NORMAL   ACTIVE        592          0     6      6
```

4. Enter **datapath query adapter** and press **Enter**. In the example shown in the following output, eight devices are attached to the SCSI path:

```
Total Devices : 6


DEV#:   0  DEVICE NAME: Disk1 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06D23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk1 Part0     OPEN   NORMAL        108          0

DEV#:   1  DEVICE NAME: Disk2 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06E23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk2 Part0     OPEN   NORMAL         96          0

DEV#:   2  DEVICE NAME: Disk3 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06F23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk3 Part0     OPEN   NORMAL         96          0

DEV#:   3  DEVICE NAME: Disk4 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07023922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk4 Part0     OPEN   NORMAL         94          0

DEV#:   4  DEVICE NAME: Disk5 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07123922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk5 Part0     OPEN   NORMAL         90          0

DEV#:   5  DEVICE NAME: Disk6 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07223922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk6 Part0     OPEN   NORMAL         98          0
```

Activating additional paths

Perform the following steps to activate additional paths to a vpath device:

1. Install any additional hardware on the Windows Server 2003 server or the ESS.
2. Click **Start → Program → Administrative Tools → Computer Management**.
3. Click **Device Manager**.
4. Right-click **Disk drives**.

5. Click **Scan for hardware changes**.
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 SDD devices, verify that the additional paths have been installed correctly.

Perform the following steps to verify that the additional paths have been installed correctly:

1. Click **Start → Program → Subsystem Device Driver → Subsystem Device Driver Management**. An MS-DOS window opens.
2. Enter **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 host bus adapter has been installed:

```
Active Adapters :2

Adpt#     Adapter Name   State     Mode     Select     Errors  Paths  Active
    0  Scsi Port4 Bus0  NORMAL   ACTIVE        592          0     6      6
    1  Scsi Port5 Bus0  NORMAL   ACTIVE        559          0     6      6
```

3. Enter **datapath query adapter** 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 host bus adapter and the new device numbers that were assigned. The following output is displayed:

```
Total Devices : 6


DEV#:   0  DEVICE NAME: Disk1 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06D23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk1 Part0     OPEN   NORMAL        108          0
    1    Scsi Port5 Bus0/Disk1 Part0     OPEN   NORMAL         96          0

DEV#:   1  DEVICE NAME: Disk2 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06E23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk2 Part0     OPEN   NORMAL         96          0
    1    Scsi Port5 Bus0/Disk2 Part0     OPEN   NORMAL         95          0

DEV#:   2  DEVICE NAME: Disk3 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 06F23922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk3 Part0     OPEN   NORMAL         96          0
    1    Scsi Port5 Bus0/Disk3 Part0     OPEN   NORMAL         94          0

DEV#:   3  DEVICE NAME: Disk4 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07023922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk4 Part0     OPEN   NORMAL         94          0
    1    Scsi Port5 Bus0/Disk4 Part0     OPEN   NORMAL         96          0

DEV#:   4  DEVICE NAME: Disk5 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07123922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk5 Part0     OPEN   NORMAL         90          0
    1    Scsi Port5 Bus0/Disk5 Part0     OPEN   NORMAL         99          0

DEV#:   5  DEVICE NAME: Disk6 Part0  TYPE: 2105800    POLICY: OPTIMIZED
SERIAL: 07223922
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port4 Bus0/Disk6 Part0     OPEN   NORMAL         98          0
    1    Scsi Port5 Bus0/Disk6 Part0     OPEN   NORMAL         79          0
```

Error recovery and retry policy

With SDD 1.5.0.0 (or later), error recovery policy is designed to report failed I/O requests to applications more quickly. This process prevents unnecessary retries, which can cause the I/O activities on good paths and vpath devices to halt for an unacceptable period of time.

SDD 1.5.0.0 (or later) error recovery policies support the following modes of operation:

single-path mode (for ESS only)

An Windows Server 2003 host system has only one path that is configured to an ESS logical unit number (LUN).

SDD, in single-path mode, has the following characteristics:

- When an I/O error occurs, SDD retries the I/O operation up to two times.

- With the SDD 1.5.0.0 (or later) error recovery policy, SDD returns the failed I/O to the application and sets the state of this failing path to DEAD. SDD driver relies on the SDD server daemon to detect the recovery of the single path. The SDD server daemon recovers this failing path and changes its state to OPEN. (SDD can change a single and failing path into DEAD state.)

multipath mode

The host system has multiple paths that are configured to a supported storage device.

SDD 1.5.0.0 (or later) error recovery policies in multipath mode have the following common characteristics:

- If an I/O error occurs on the last operational path to a device, SDD attempts to reuse (performs a failback operation to return to) a previously failed path.

The SDD 1.5.0.0 (or later) error recovery policy in multipath mode has the following latest characteristics:

- SDD 1.5.0.0 (or later) does not attempt to use the path until three successful I/O operations occur on an operational path.
- If an I/O error occurs consecutively on a path and the I/O error count reaches three, SDD immediately changes the state of the failing path to DEAD.
- Both SDD driver and the SDD server daemon can put a last path into DEAD state, if this path is no longer functional. The SDD server can automatically change the state of this path to OPEN after it is recovered. Alternatively, you can manually change the state of the path to OPEN after it is recovered by using the **datapath set path online** command. Go to “datapath set device path” on page 304 for more information.
- If an I/O fails on all OPEN paths to an ESS LUN, SDD returns the failed I/O to the application and changes the state of all OPEN paths (for failed I/Os) to DEAD, even if some paths did not reach an I/O error count of three.
- If an OPEN path already failed some I/Os, it will not be selected as a retry path.

Note: When a path failover does not work with the QLogic card, you need to verify that the Target Enabled bit is set in the QLogic BIOS.

Press **Ctrl+Q** at boot time to change the QLogic BIOS settings.

Refer to the following URL for recommended QLogic BIOS settings:

http://publibfp.boulder.ibm.com/epubs/pdf/f2bhs00.pdf

Support for Windows Server 2003 clustering

SDD 1.5.0.0 (or later) is required to support Windows Server 2003 clustering. SDD 1.5.0.0 (or later) does not support I/O load balancing in a Windows Server 2003 clustering environment.

When running Windows Server 2003 clustering, clustering failover might not occur when the last path is being removed from the shared resources. See *Microsoft article Q294173* for additional information. Windows Server 2003 does not support dynamic disks in the MSCS environment.

Special considerations in the Windows Server 2003 clustering environment

There are subtle differences in the way that SDD handles path reclamation in a Windows Server 2003 clustering environment compared to a nonclustering environment. When the Windows Server 2003 server loses a path in a nonclustering environment, the path condition changes from open to dead and the adapter condition changes from active to degraded. The adapter and path condition will not change until the path is made operational again. When the Windows Server 2003 server loses a path in a clustering environment, the path condition changes from open to dead and the adapter condition changes from active to degraded. However, after a period of time, the path condition changes back to open and the adapter condition changes back to normal, even if the path has not been made operational again.

Note: The adapter goes to DEGRAD state when there are active paths left on the adapter. It goes to FAILED state when there are no active paths.

The **datapath set adapter # offline** command operates differently in a clustering environment as compared to a nonclustering environment. In a clustering environment, the **datapath set adapter offline** command does not change the condition of the path if the path is active or being reserved. If you issue the command, the following message is displayed: `to preserve access some paths left online`.

Configure Windows 2003 cluster with SDD installed

The following variables are used in this procedure:

server_1	Represents the first server with two host bus adapters (HBAs).
server_2	Represents the second server with two HBAs.
hba_a	Represents the first HBA for *server_1*.
hba_b	Represents the second HBA for *server_1*.
hba_c	Represents the first HBA for *server_2*.
hba_d	Represents the second HBA for *server_2*.

Perform the following steps to configure a Windows Server 2003 cluster with SDD:

1. Configure LUNs on the ESS as shared for all HBAs on both *server_1* and *server_2*.
2. Connect *hba_a* to the ESS, and restart *server_1*.
3. Click **Start → Programs → Administrative Tools → Computer Management**. The Computer Management window opens. From the Computer Management window, select Storage and then Disk Management to work with the storage devices attached to the host system.

   The operating system will recognize each additional path to the same LUN as a device.
4. Disconnect *hba_a* and connect *hba_b* to the ESS. Restart *server_1*.

5. Click **Start → Programs → Administrative Tools → Computer Management**. The Computer Management window opens. From the Computer Management window, select Storage and then Disk Management to verify the correct number of LUNs that are connected to *server_1*.

   If the number of LUNs that are connected to *server_1* is correct, proceed to 6.

   If the number of LUNs that are connected to *server_1* is incorrect, perform the following steps:

   a. Verify that the cable for *hba_b* is connected to the ESS.
   b. Verify your LUN configuration on the ESS.
   c. Repeat steps 2 - 5.
6. Install SDD on *server_1*, and restart *server_1*.

   For installation instructions, go to “Installing SDD 1.5.0.0 (or later)” on page 248 section.
7. Connect *hba_c* to the ESS, and restart *server_2*.
8. Click **Start → Programs → Administrative Tools → Computer Management**. The Computer Management window opens. From the Computer Management window, select Storage and then Disk Management to verify the correct number of LUNs that are connected to *server_2*.

   The operating system will see each additional path to the same LUN as a device.
9. Disconnect *hba_c* and connect *hba_d* to the ESS. Restart *server_2*.

10. Click **Start → Programs → Administrative Tools → Computer Management**. The Computer Management window is displayed. From the Computer Management window, select Storage and then Disk Management to verify the correct number of LUNs that are connected to *server_2*.

    If the number of LUNs that are connected to *server_2* is correct, proceed to 11.

    If the number of LUNs that are connected to *server_2* is incorrect, perform the following steps:

    a. Verify that the cable for *hba_d* is connected to the ESS.

    b. Verify your LUN configuration on the ESS.

    c. Repeat steps 7 - 10.

11. Install SDD on *server_2*, and restart *server_2*.

    For installation instructions, go to “Installing SDD 1.5.0.0 (or later)” on page 248.

12. Connect both *hba_c* and *hba_d* on *server_2* to the ESS, and restart *server_2*.
13. Use the **datapath query adapter** and **datapath query device** commands to verify the correct number of LUNs and paths on *server_2*.
14. Click **Start → Programs → Administrative Tools → Computer Management**. The Computer Management window opens. From the Computer Management window, select Storage and then Disk Management to verify that the actual number of LUNs as online devices is correct.
15. Format the raw devices with NTFS.

    Make sure to keep track of the assigned drive letters on *server_2*.

16. Connect both *hba_a* and *hba_b* on *server_1* to the ESS, and restart *server_1*.
17. Use the **datapath query adapter** and **datapath query device** commands to verify the correct number of LUNs and paths on *server_1*.

    Verify that the assigned drive letters on *server_1* match the assigned drive letters on *server_2*.

18. Restart *server_2*.
    - Install the MSCS software on *server_1*, restart *server_1*, reapply Service Pack 2 or higher to *server_1*, and restart *server_1* again.
    - Install the MSCS software on *server_2*, restart *server_2*, reapply Service Pack 2 to *server_2*, and restart *server_2* again.
19. Use the **datapath query adapter** and **datapath query device** commands to verify 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.

Information about installing a Windows 2003 cluster can be found in a file, confclus.exe, located at:

www.microsoft.com/downloads/details.aspx?displaylang=en&familyid=96F76ED7-9634-4300-9159-89638F4B4EF7

Chapter 11. Using the SDD server and the SDDPCM server

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 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 chapter in this guide for the specific platform for modifications to SDD server daemon operation.

Note: sddsrv 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 executions unless otherwise specified for a specific platform. If the test succeeds, then sddsrv reclaims these paths and changes the states of these paths according to the following characteristics:

- If the state of the vpath is OPEN, then sddsrv changes the states of INVALID and CLOSE_DEAD paths of that vpath to OPEN.
- If the state of the vpath is CLOSE, then sddsrv changes the states of CLOSE_DEAD paths of that vpath to CLOSE.
- Sddsrv 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 executions 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 vpath is in the OPEN state and the path is not working, then sddsrv changes the state of the path from OPEN to DEAD.
- If the vpath is in the CLOSE state and the path is not working, then sddsrv changes the state of the path from CLOSE to CLOSE_DEAD.
- sddsrv will put the last path to DEAD or CLOSE_DEAD depending upon the state of vpath.

Note: Sddsrv will not test paths that are manually placed offline.

In SDD 1.5.0.x (or earlier), sddsrv by default was binding to a TCP/IP port and listening for incoming requests. In SDD 1.5.1.x (or later), sddsrv does not bind to

© Copyright IBM Corp. 1999, 2004

any TCP/IP port by default, but allows port binding to be dynamically enabled or disabled. For all platform except Linux, the SDD package ships a template file of sddsrv.conf that is named sample_sddsrv.conf. On all UNIX platforms except Linux, the sample_sddsrv.conf file is located in the /etc directory. On Windows platforms, the sample_sddsrv.conf file is in the directory where SDD is installed. You must use the sample_sddsrv.conf file to create the sddsrv.conf file in the same directory as sample_sddsrv.conf by simply copying it and naming the copied file sddsrv.conf. You can then dynamically change port binding by modifying parameters in sddsrv.conf.

Because sddsrv's TCP/IP interface is disabled by default, you cannot get sddsrv 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 sddsrv trace is saved in sddsrv.log and sddsrv_bak.log files. The sddsrv trace log files are wrap-around files and each file is maximum of 4 MB in size. sddsrv also collects the SDD driver trace and puts it in log files. It creates sdd.log and sdd_bak.log files for the driver trace. The SDD driver trace log files are also wrap-around files and each file is maximum of 4 MB in size. You will find sddsrv.log, sddsrv_bak.log, sdd.log and sdd_bak.log files in the following directory based on your host system platform:

- AIX - /var/adm/ras
- HP-UX - /var/adm
- Linux - /var/log
- Solaris- /var/adm
- Windows 2000 and Windows NT - \WINNT\system32
- Windows Server 2003 - \Windows\system32

See "SDD data collection for problem analysis," on page 307 for information about reporting SDD problems.

sddsrv and the IBM TotalStorage Expert V.2.1.0

The IBM TotalStorage Expert V.2.1.0 needs to communicate with sddsrv through a TCP/IP socket on the port on which sddsrv is running. The sddsrv 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. You should apply your corporate security rules to this port.

sddsrv and IBM TotalStorage support for Geographically Dispersed Sites for Microsoft Cluster Service

The sddsrv 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). You should 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. It tests INVALID paths, and it tests OPEN_FAILED paths when healthcheck is turned off. It also tests CLOSE_FAILED paths for devices that are in the CLOSED state. The daemon "sleeps" for one-minute intervals between consecutive executions. If the test succeeds, then pcmsrv reclaims these paths and changes the states of these paths according to the following characteristics:

- If the state of the device is OPEN, then pcmsrv changes the states of INVALID paths of that device to OPEN.
- If the state of the device is OPEN, and healthcheck function is turned off, then pcmsrv changes the states of OPEN_FAILED paths of that device to OPEN.
- If the state of the device is CLOSE, then pcmsrv changes the states of CLOSE_FAILED paths of the device to CLOSE.

pcmsrv 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 simply 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 pcmsrv is saved in pcmsrv.log and pcmsrv_bak.log files. These are wrap-around files and each is a maximum of 4 MB in size. Trace files are located in the /var/adm/ras directory.

sddsrv.conf and pcmsrv.conf file format

The sddsrv.conf and pcmsrv.conf files contain the following parameters:

- enableport - This parameter allows you to enable or disable sddsrv or pcmsrv 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 or pcmsrv.
- loopbackbind - If you set the enableport parameter to true, then the loopbackbind parameter specifies whether sddsrv or pcmsrv will listen to any Internet address or the loopback (127.0.0.1) address. To enable sddsrv or pcmsrv to listen to any Internet address, the loopbackbind parameter must be set to false. To enable sddsrv or pcmsrv to listen only to the loopback address 127.0.0.1, the loopbackbind parameter must be set to true.
- portnumber - This parameter specifies the port number that sddsrv or pcmsrv will bind 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, then this parameter must be set to a valid port number to which sddsrv or pcmsrv can bind. Use a port number that is not used by any other application.

You can modify these parameters while sddsrv or pcmsrv is executing to enable or disable the TCP/IP interface 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 will enable sddsrv or pcmsrv to bind to a TCP/IP port.

enableport = false will disable 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*, then only sddsrv or pcmsrv will bind to the port number specified. The default value of this parameter is 20001.

Chapter 12. Using the datapath commands

SDD provides commands that you can use to:

- 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 suspend or resume adapters.
- 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.

This chapter includes descriptions of these commands. Table 33 provides an alphabetical list of these commands, a brief description, and where to go in this chapter for more information.

Table 33. Commands

Command	Description	Page
datapath disable ports	Places paths connected to certain ports offline.	284
datapath enable ports	Places paths connected to certain ports online.	285
datapath open device path	Dynamically opens a path that is in an Invalid or Close_Dead state.	286
datapath query adapter	Displays information about adapters.	288
datapath query adaptstats	Displays performance information for all SCSI and FCS adapters that are attached to SDD devices.	290
datapath query device	Displays information about devices.	291
datapath query devstats	Displays performance information for a single SDD device or all SDD devices.	293
datapath query essmap	Displays each device, path, location, and attributes.	295
datapath query portmap	Displays ESS vpath, device, port, and location.	296
datapath query wwpn	Displays the World Wide Port Name (WWPN) of the host fibre-channel adapters.	297
datapath remove adapter	Dynamically removes an adapter.	298
datapath remove device path	Dynamically removes a path of a device.	299
datapath set adapter	Sets all device paths that are attached to an adapter to online or offline.	301
datapath set device policy	Dynamically changes the path-selection policy of the SDD device or devices.	303
datapath set device path	Sets the path of a device to online or offline.	304
datapath set qdepth	Dynamically enables or disables queue depth.	305

© Copyright IBM Corp. 1999, 2004

datapath disable ports

The **datapath disable ports** command sets SDD device paths offline for specified ESS location code.

Note: This command is supported for AIX host systems only.

Syntax

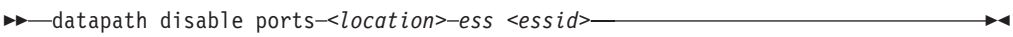

Parameters

location
The location 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

essid
The ESS serial number, given by the output of **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  ALGORITHM:  Load Balance
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 device paths online for specified ESS location code.

Note: This command is supported for AIX host systems only.

Syntax

```
►►—datapath enable ports—location—ess essid————————————————————————►◄
```

Parameters

location
The location 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

essid
The ESS serial number, given by the output of **datapath 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  ALGORITHM:  Load Balance
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 Invalid or Close_Dead state. You can use this command even when the I/O is actively running.

Note: This command is supported for Sun, HP, and AIX 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 issuing 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. The policy selected is one of the following policies: Optimized (another name for load-balancing), Round Robin, and Failover only.

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.
Suspend
Path is suspended and is not being used.

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—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 adapter** command, the following output is displayed:

```
Active Adapters :4

Adpt#     Adapter Name   State     Mode      Select     Errors  Paths  Active
    0            scsi3  NORMAL   ACTIVE  129062051          0     64       0
    1            scsi2  NORMAL   ACTIVE   88765386        303     64       0
    2           fscsi2  NORMAL   ACTIVE  407075697       5427   1024       0
    3           fscsi0  NORMAL   ACTIVE  341204788      63835    256       0
```

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.

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.
Suspend
The adapter is currently suspended. See “datapath set adapter” on page 301

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.

Note: Windows 2000 and Windows Server 2003 host systems can display different values for State and Mode depending on adapter type when a path is placed offline due to a bay quiescence.

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 or all 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 a specific device number.

Syntax

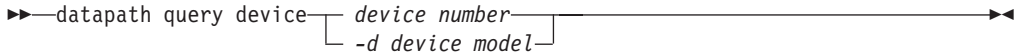

Parameters

device number
The device number refers to the device *index* number as displayed by the **datapath query device** command, rather than the SDD device number.

-d device model [1]
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:

2105 Display all 2105 models (ESS).

2105F Display all 2105 F models (ESS).

2105800
Display all 2105 800 models (ESS).

2145 Display all 2145 models (SAN Volume Controller).

2062 Display all 2062 models (SAN Volume Controller for Cisco MDS 9000).

Examples

If you enter the **datapath query device 0** command, the following output is displayed:

For ESS:

```
DEV#: 0   DEVICE NAME: vpath0    TYPE: 2105E20          POLICY: Optimized
SERIAL: 31412028
============================================================================
Path#     Adapter/Hard Disk              State     Mode    Select    Errors
0          fscsi0/hdisk2                  OPEN    NORMAL         9         0
1          fscsi0/hdisk4                  OPEN    NORMAL        12         0
2          fscsi1/hdisk6                  OPEN    NORMAL        21         0
3          fscsi1/hdisk8                  OPEN    NORMAL        23         0
```

For SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000:

```
DEV#:   0  DEVICE NAME: vpath7  TYPE: 2145           POLICY:    Optimized
 SERIAL: 6005676801800210B000000000000007
 ============================================================================
 Path#      Adapter/Hard Disk          State     Mode     Select     Errors
     0          fscsi0/hdisk9          CLOSE   NORMAL        492          0
```

1. The option to specify a device model cannot be used when you query a specific device number.

```
    1         fscsi0/hdisk18          CLOSE   NORMAL          0          0
    2         fscsi1/hdisk27          CLOSE   NORMAL        541          0
    3         fscsi1/hdisk36          CLOSE   NORMAL          0          0
```

Notes:

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, you should always run the **datapath query device** command first.
2. For SDD 1.4.0.0 (or later), the location of Policy and Serial Number are swapped.

The terms used in the output are defined as follows:

Dev# The number of this device defined by SDD.

Name The name of this device defined by SDD.

Type The device product ID from inquiry data.

Policy The current path selection policy selected for the device. The policy selected is one of the following policies: Optimized (another name for load-balancing), Round Robin, and Failover only.

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.
Suspend
The path is currently suspended.

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 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 a specific device number.

Syntax

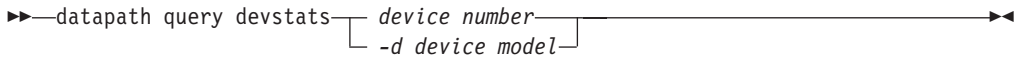

Parameters

device number
The device number refers to the device *index* number as displayed by the **datapath query device** command, rather than the SDD device number.

-d device model [2]
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:

2105 Display all 2105 models (ESS).

2105F Display all 2105 F models (ESS).

2105800
Display all 2105 800 models (ESS).

2145 Display all 2145 models (SAN Volume Controller).

2062 Display all 2062 models (SAN Volume Controller for Cisco MDS 9000).

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

2. The option to specify a device model cannot be used when you query a specific device number.

- 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: This command is supported for AIX host systems only.

Syntax

```
►►—datapath query essmap————————————————————————————————————►◄
```

Examples

If you enter the **datapath query essmap** command, the following output is displayed:

```
 Disk      Path        Location     adapter   LUN SN       Type             Size   LSS    Vol   Rank  C/A   S   Connection    port  RaidMode
-------   ------      ---------    --------  -------      ----             ----   ---    ---   ----  ---   -   ----------    ----  --------
vpath0     hdisk2      10-68-01[FC] fscsi0   20112028  IBM 2105-E20          1.0    18     1    1201   02   Y   R1-B1-H3-ZA     8          5
vpath0     hdisk3      10-68-01[FC] fscsi0   20112028  IBM 2105-E20          1.0    18     1    1201   02   Y   R1-B3-H4-ZA    8c          5
vpath0     hdisk4      20-58-01[FC] fscsi1   20112028  IBM 2105-E20          1.0    18     1    1201   02   Y   R1-B1-H3-ZA     8          5
vpath0     hdisk5      20-58-01[FC] fscsi1   20112028  IBM 2105-E20          1.0    18     1    1201   02   Y   R1-B3-H4-ZA    8c          5
vpath1     hdisk6      10-68-01[FC] fscsi0   20712028  IBM 2105-E20          1.0    18     7    1201   02   Y   R1-B1-H3-ZA     8          5
vpath1     hdisk7      10-68-01[FC] fscsi0   20712028  IBM 2105-E20          1.0    18     7    1201   02   Y   R1-B3-H4-ZA    8c          5
vpath1     hdisk8      20-58-01[FC] fscsi1   20712028  IBM 2105-E20          1.0    18     7    1201   02   Y   R1-B1-H3-ZA     8          5
vpath1     hdisk9      20-58-01[FC] fscsi1   20712028  IBM 2105-E20          1.0    18     7    1201   02   Y   R1-B3-H4-ZA    8c          5
```

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 device.

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 ESS.

Type The device and model.

Size The configured capacity of the LUN in gigabytes.

LSS The logical subsystem where the LUN resides.

Vol The volume number within the ESS.

Rank The unique identifier for each RAID-array within the ESS.

C/A The cluster and adapter accessing the array.

S Indicates that the device is shared by two and more ESS ports. Valid values are *yes* or *no*.

Connection The physical location code of ESS adapter through which the LUN is accessed.

Port The ESS port through which the LUN is accessed.

RaidMode The disk RAID mode.

datapath query portmap

The **datapath query portmap** command displays hdisk status from an ESS physical location code perspective.

Note: This command is supported for AIX host systems only.

Syntax

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
                  A B A B A B A B    A B A B A B A B    A B A B A B A B    A B A B A B A B
26317    vpath10   - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
26317    vpath11   - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
26317    vpath12   - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
26317    vpath13   - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
26317    vpath14   - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
26317    vpath15   - - - - - - - -    Y - - - - - - -    - - - - - - - -    Y - - - - - - -
26317    vpath16   - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
26317    vpath17   - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -
12028    vpath18   - - - - - - - -    O - - - - - - -    - - - - - - - -    O - - - - - - -

Y  =  online/open
O  =  online/closed
N  =  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.

O The port is online and closed, meaning that at least one path state/mode is closed/online.

N 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.

datapath query wwpn

The **datapath query wwpn** command displays the World Wide Port Name (WWPN) of the host fibre-channel adapters.

Note: This command is available for AIX, Linux, Windows NT, and Windows 2000 host systems only.

Syntax

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 will fail.

Note: The newly installed adapter can be reclaimed with the **addpaths** command. For detailed procedures, see “addpaths” on page 85.

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#     Adapter 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 entry for Adpt# 0 disappears from the **datapath query adapter** command output.
- All paths attached to adapter 0 disappear from the **datapath query device** command output.
- This can be done while I/O is running.

```
+-----------------------------------------------------------------------------+
|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         27     10      10|
|    3           fscsi3  NORMAL   ACTIVE       4997         35     10      10|
+-----------------------------------------------------------------------------+
```

Note that Adpt# 0 fscsi0 is removed and the Select counts are increased.

datapath remove device path

The **datapath remove device path** command dynamically removes a specific path of a SDD device that is attached to an adapter. SDD reserves the last path of a device. If the requested path is the only accessible path to a device, the command will fail.

Note: A removed path can be reclaimed by using the **addpaths** command. See "addpaths" on page 85.

Note: 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#'s 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|
+-----------------------------------------------------------------+
```

Note that 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 device paths attached to an adapter to:

- Online
- Offline
- Suspend (AIX host systems only)
- Resume (AIX host systems only)

SDD reserves the last path of a device. If all accessible paths to a device are attached to one adapter, the offline or suspend option will fail.

Note: The **datapath set adapter offline** command can be entered even when the device is closed.

Syntax

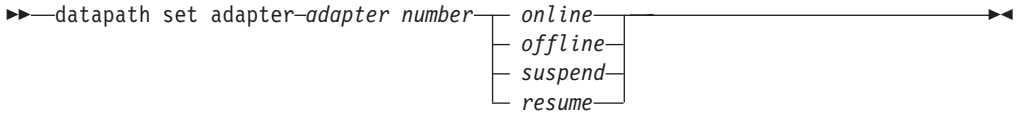

Parameters

adapter number
The index number of the adapter that you want to change.

online
Sets the adapter online.

offline
Sets the adapter offline.

suspend
Suspends the adapter. (AIX host systems only)

resume
Resumes the adapter. (AIX host systems only)

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 offlined adapter 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 device state.

If you enter the **datapath set adapter 0 suspend** command:

- Adapter 0's mode changes to OFFLINE and its state to SUSPEND.

- The mode of all paths attached to adapter 0 changes to OFFLINE and their state to SUSPEND.
- All paths attached to this suspended adapter are closed.

A suspended adapter can be *resumed* by issuing the **datapath set adapter 0 resume** command:

- 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 device state.

datapath set device policy

The **datapath set device policy** command dynamically changes the path selection policy per SDD vpath device.

Syntax

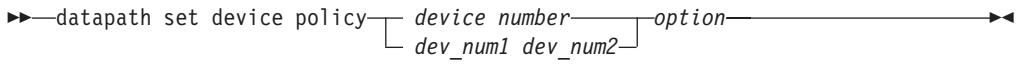

Note: This command is available for AIX, Linux, HP, and SUN host systems only.

Parameters

device number
The device number refers to the device index number, rather than the SDD device number.

device number1 device number2
When 2 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:

- **rr**, where *rr* indicates round robin
- **lb**, where *lb* indicates load balancing (also known as optimized policy)
- **df**, where *df* indicates the default policy, which is load balancing
- **fo**, where *fo* indicates failover policy

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

Examples

If you enter **datapath set device 2 7 policy rr**, the path-selection policy of 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 a device either to online or offline. You cannot remove the last path to a device from service. This prevents a data access failure from occurring. The **datapath set device path** command can be entered even when the device is closed.

Syntax

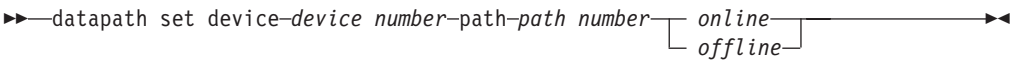

Parameters

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 qdepth

The **datapath set qdepth** command dynamically sets the queue depth logic to either enable or disable even when I/O is running.

Note: This command is available for AIX host systems only.

Syntax

Parameters

Enable
Enable the queue depth logic.

Disable
Disable the queue depth logic.

Examples

If you enter the **datapath set qdepth disable** command, the following output is displayed:

```
+-----------------------------------------------------------------------+
|Success: set qdepth_enable to no                                       |
+-----------------------------------------------------------------------+
```

The SDD ODM attribute, qdepth_enable, is updated. The following output is displayed when you enter **lsattr -El dpo**.

```
+-----------------------------------------------------------------------+
|persistent_resv yes Subsystem Supports Persistent Reserve Command False|
|qdepth_enable   no  Queue Depth Control                          True  |
+-----------------------------------------------------------------------+
```

Appendix. SDD data collection for problem analysis

Starting with SDD 1.5.1.0 (and later), SDD has enhanced its trace capability to help with 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

These files can be found in the following directories:

- AIX - /var/adm/ras
- HP-UX - /var/adm
- Linux - /var/log
- Solaris - /var/adm
- Windows 2000 and Windows NT - \WINNT\system32
- Windows Server 2003 - \Windows\system32

When you report a SDD problem, the following information is essential for problem determination:

- sdd.log
- sdd_bak.log
- sddsrv.log
- sddsrv_bak.log
- Output of the **datapath query adapter** command
- Output of the **datapath query device** command
- System error log

© Copyright IBM Corp. 1999, 2004

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.

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.

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.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

© Copyright IBM Corp. 1999, 2004

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 License Agreement for Non-Warranted Programs.

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.

Trademarks

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

AIX
AS/400
DFSMS/MVS
Enterprise Storage Server
The eServer logo
ES/9000
ESCON
FICON
FlashCopy
GDPS
HACMP/6000
IBM
IBM logo
iSeries
MVS
MVS/ESA
Netfinity
NetVista
NUMA-Q
Operating System/400
OS/390
OS/400
pSeries
RS/6000
S/390
Seascape

SNAPSHOT
SP
System/360
System/370
System/390
System/400
TotalStorage
Versatile Storage Server
VM/ESA
VSE/ESA
xSeries
zSeries
z/Architecture
z/OS

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

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

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

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

IBM agreement for licensed internal 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. SEE “IBM agreement for licensed internal code.”

You accept the terms of this Agreement[3] by your initial use of a machine that contains IBM Licensed Internal Code (called “Code”). These terms apply to Code used by certain machines IBM or your reseller specifies (called “Specific Machines”). International Business Machines Corporation or one of its subsidiaries (“IBM”) owns copyrights in Code or has the right to license Code. IBM or a third party owns all copies of Code, including all copies made from them.

If you are the rightful possessor of a Specific Machine, IBM grants you a license to use the Code (or any replacement IBM provides) on, or in conjunction with, only the Specific Machine for which the Code is provided. IBM licenses the Code to only one rightful possessor at a time.

Under each license, IBM authorizes you to do only the following:

3. Form Z125-4144

1. execute the Code to enable the Specific Machine to function according to its Official Published Specifications (called “Specifications”);
2. make a backup or archival copy of the Code (unless IBM makes one available for your use), provided you reproduce the copyright notice and any other legend of ownership on the copy. You may use the copy only to replace the original, when necessary; and
3. execute and display the Code as necessary to maintain the Specific Machine.

You agree to acquire any replacement for, or additional copy of, Code directly from IBM in accordance with IBM’s standard policies and practices. You also agree to use that Code under these terms.

You may transfer possession of the Code to another party only with the transfer of the Specific Machine. If you do so, you must 1) destroy all your copies of the Code that were not provided by IBM, 2) either give the other party all your IBM-provided copies of the Code or destroy them, and 3) notify the other party of these terms. IBM licenses the other party when it accepts these terms. These terms apply to all Code you acquire from any source.

Your license terminates when you no longer rightfully possess the Specific Machine.

Actions you must not take

You agree to use the Code only as authorized above. You must not do, for example, any of the following:

1. Otherwise copy, display, transfer, adapt, modify, or distribute the Code (electronically or otherwise), except as IBM may authorize in the Specific Machine’s Specifications or in writing to you;
2. Reverse assemble, reverse compile, or otherwise translate the Code unless expressly permitted by applicable law without the possibility of contractual waiver;
3. Sublicense or assign the license for the Code; or
4. Lease the Code or any copy of it.

Glossary

This glossary includes terms for the IBM TotalStorage Enterprise Storage Server (ESS) and other Seascape solution 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.
- The *IBM Glossary of Computing Terms* is available online at the following Web site: http://www.ibm.com/ibm/terminology/. 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.

Numerics

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*.

3390. The machine number of an IBM disk storage system. The ESS, when interfaced to IBM S/390 or zSeries hosts, is set up to appear as one or more 3390 devices, with a choice of 3390-2, 3390-3, or 3390-9 track formats.

3990. The machine number of an IBM control unit.

7133. The machine number of an IBM disk storage system. The Model D40 and 020 drawers of the 7133 can be installed in the 2105-100 expansion enclosure of the ESS.

8-pack. See *disk eight pack*.

A

access. (1) To obtain the use of a computer resource. (2) In computer security, a specific type of interaction between a subject and an object that results in flow of information from one to the other.

access-any mode. One of the two access modes that can be set for the ESS during initial configuration. It enables all fibre-channel-attached host systems with no defined access profile to access all logical volumes on the ESS. 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*.

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*.

© Copyright IBM Corp. 1999, 2004

allocated storage. In an ESS, 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 ESS, 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 an ESS, 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 ESS, 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. (1) Either of the digits 0 or 1 when used in the binary numeration system. (T) (2) The storage medium required to store a single binary digit. See also *byte*.

block. (1) A string of data elements recorded or transmitted as a unit. The elements may be characters, words, or physical records. (T) (2) In the ESS, a group of consecutive bytes used as the basic storage unit in fixed-block architecture (FBA). All blocks on the storage device are the same size (fixed size). See also *fixed-block architecture* and *data record*.

byte. (1) A group of eight adjacent binary digits that represent one EBCDIC character. (2) The storage medium required to store eight bits. See also *bit*.

C

cache. A special-purpose buffer storage, smaller and faster than main storage, used to hold a copy of instructions and data obtained from main storage and likely to be needed next by the processor. (T)

cache fast write. In the ESS, 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 ESS and a service provider. The ESS can use this link to place a call to IBM or to another service provider

when it requires service. With access to the machine, service personnel can perform service tasks, such as viewing error logs and problem logs or initiating trace and dump retrievals. See also *heartbeat* and *remote technical assistance information network*.

cascading. (1) Connecting network controllers to each other in a succession of levels, to concentrate many more lines than a single level permits. (2) In high-availability cluster multiprocessing (HACMP), cascading pertains to a cluster configuration in which the cluster node with the highest priority for a particular resource acquires the resource if the primary node fails. The cluster node relinquishes the resource to the primary node upon reintegration of the primary node into the cluster.

catcher. A server that service personnel use to collect and retain status data that an ESS 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 ESS 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. (1) In the ESS, a partition capable of performing all ESS functions. With two clusters in the ESS, any operational cluster can take over the processing of a failing cluster. (2) In the AIX operating system, a group of nodes within a complex.

cluster processor complex (CPC). In the ESS, the unit within a cluster that provides the management function for the ESS. It consists of cluster processors, cluster memory, and related logic.

command-line interface (CLI). An interface provided by an operating system that defines a set of commands and enables a user (or a script-like language) to issue these commands by typing text in response to the command prompt (for example, DOS commands, UNIX shell commands). See also *Copy Services command-line interface*.

compression. (1) The process of eliminating gaps, empty fields, redundancies, and unnecessary data to shorten the length of records or blocks. (2) Any encoding that reduces the number of bits used to represent a given message or record. (GC)

computer-electronic complex (CEC). The set of hardware facilities associated with a host computer.

concurrent copy. A facility on a storage server that enables a program to make a backup of a data set while the logical volume remains available for subsequent processing. The data in the backup copy is frozen at the point in time that the server responds to the request.

concurrent installation of licensed internal code. Process of installing licensed internal code on an ESS 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). (1) A device that coordinates and controls the operation of one or more input/output devices, and synchronizes the operation of such devices with the operation of the system as a whole. (2) 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. (3) 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 zSeries 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 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*.

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 zSeries 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). (1) A mass storage medium on which a computer stores data. (2) A disk device.

disk cage. A container for disk drives. Each disk cage supports eight disk 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*.

distributed file service (DFS). A service that provides data access over IP networks.

DNS. See *domain name system*.

domain. (1) That part of a computer network in which the data processing resources are under common control. (2) In TCP/IP, the naming system used in hierarchical networks. (3) A Copy Services server group, in other words, the set of clusters 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. (1) A peripheral device, especially one that has addressed storage media. See also *disk drive module*. (2) The mechanism used to seek, read, and write information on a storage medium.

duplex. (1) Regarding ESS Copy Services, the state of a volume pair after PPRC has completed the copy operation and the volume pair is synchronized. (2) In general, pertaining to a communication mode in which data can be sent and received at the same time.

dynamic sparing. The ability of a storage server to move data from a failing disk drive module (DDM) to a spare DDM while maintaining storage functions.

E

E10. The predecessor of the F10 model of the 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). (1) An Enterprise Systems Architecture/390 and zSeries computer peripheral interface. The I/O interface uses ESA/390 logical protocols over a serial interface that configures attached units to a communication fabric. (2) A set of IBM products and services that provide a dynamically connected environment within an enterprise.

EPO. See *emergency power off*.

ERDS. See *error-recording data set*.

error-recording data set (ERDS). On S/390 and zSeries hosts, a data set that records data-storage and data-retrieval errors. A service information message (SIM) provides the error information for the ERDS.

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 zSeries 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 zSeries 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 @server*.

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 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. Refer to American National Standards Institute (ANSI) X3T11/93-275. In this topology, two or more fibre-channel end points are interconnected through a looped interface. 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 SBCON for ESA/390 command protocols). Refer to 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 10^9 bytes. A gigabyte of memory is 2^{30} 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 Multi-Processing*.

hard disk drive (HDD). (1) A storage medium within a storage server used to maintain information that the storage server requires. (2) A mass storage medium for computers that is typically available as a fixed disk (such as the disks used in system units of personal computers or in drives that are external to a personal computer) or a removable cartridge.

hardware service manager (HSM). An option on an AS/400 or iSeries host that enables the user to display and work with system hardware resources and to debug input-output processors (IOP), input-output adapters (IOA), and devices.

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. (1) 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. (2) In AS/400 storage management, an automatic method to manage and distribute data between the different storage layers, such as disk units and tape library devices.

High-Availability Cluster Multi-Processing (HACMP). Software that provides host clustering, so that a failure of one host is recovered by moving jobs to other hosts within the cluster.

high-speed link (HSL). A hardware connectivity architecture that links system processors to system input/output buses and other system units.

home address (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 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 zSeries) or of the open-systems type, that is connected to the ESS. S/390 or zSeries 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 @server. The IBM brand name for a series of server products that are optimized for e-commerce. The products include the iSeries, pSeries, xSeries, and zSeries.

IBM product engineering (PE). The third-level of IBM service support. Product engineering is composed of IBM engineers who have experience in supporting a product or who are knowledgeable about the product.

IBM TotalStorage. The brand name used to identify storage products from IBM, including the IBM TotalStorage Enterprise Storage Server (ESS). See also *IBM TotalStorage Enterprise Storage Server* and *IBM TotalStorage Enterprise Storage Server Specialist*.

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.

IBM TotalStorage Enterprise Storage Server Subsystem Device Driver (SDD). Software that is designed to support the multipath configuration environments in the ESS. SDD resides in a host system with the native disk device driver.

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). Refer to ANSI X3.230-199x. The IP acronym is the IP in TCP/IP. See also *Transmission Control Protocol/Internet Protocol*.

invalidate. To remove a logical data unit from cache memory because it cannot support continued access to the logical data unit on the device. This removal might be the result of a failure within the storage server or a storage device that is associated with the device.

I/O. See *input/output*.

I/O adapter (IOA). 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*.

iSeries. An IBM @server product that emphasizes integration. It is the successor to the AS/400 family of servers.

J

Java Virtual Machine (JVM). A software implementation of a central processing unit (CPU) that runs compiled Java code (applets and applications). (GC)

JVM. See *Java Virtual Machine*.

K

KB. See *kilobyte*.

key field. The second (optional) field of a 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). (1) For processor storage, real, and virtual storage, and channel volume, 2^{10} or 1024 bytes. (2) For disk storage capacity and communications volume, 1000 bytes.

Korn shell. Interactive command interpreter and a command programming language.

KPOH. See *thousands of power-on hours*.

L

LAN. See *local area network*.

last-in first-out (LIFO). A queuing technique in which the next item to be retrieved is the item most recently placed in the queue. (A)

LBA. See *logical block address*.

LCU. See *logical control unit*.

least recently used (LRU). (1) The algorithm used to identify and make available the cache space that contains the least-recently used data. (2) A policy for a caching algorithm that chooses to remove from cache the item that has the longest elapsed time since its last access.

LED. See *light-emitting diode*.

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.

LIFO. See *last-in first-out*.

light-emitting diode (LED). A semiconductor chip that gives off visible or infrared light when activated.

link address. On an ESCON 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 zSeries hosts, a logical subsystem represents a logical control unit (LCU). Each control-unit image is associated with only one logical subsystem. See also *control-unit image*.

logical unit. In open systems, a logical disk drive.

logical unit number (LUN). In the SCSI protocol, a unique number 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). (1) A method of error-checking during data transfer that involves checking parity on a row of binary digits that are members of a set that forms a matrix. Longitudinal redundancy check is also called a longitudinal parity check. (2) 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. (T)

maintenance analysis procedure (MAP). A hardware maintenance document that gives an IBM service representative a step-by-step procedure for tracing a symptom to the cause of a failure.

Management Information Base (MIB). (1) A collection of objects that can be accessed by means of a network management protocol. (GC) (2) 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*.

Master Console. See *IBM TotalStorage ESS Master Console*.

MB. See *megabyte*.

MCA. See *Micro Channel architecture*.

mean time between failures (MTBF). (1) A projection of the time that an individual unit remains functional. The time is based on averaging the performance, or projected performance, of a population of statistically independent units. The units operate under a set of conditions or assumptions. (2) For a stated period in the life of a functional unit, the mean value of the lengths of time between consecutive failures under stated conditions. (I) (A)

medium. For a storage facility, the disk surface on which data is stored.

megabyte (MB). (1) For processor storage, real and virtual storage, and channel volume, 2^{20} or 1 048 576 bytes. (2) For disk storage capacity and communications volume, 1 000 000 bytes.

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.

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.

node. The unit that is connected in a fibre-channel network. An ESS is a node in a fibre-channel network.

non-RAID. A disk drive set up independently of other disk drives and not set up as part of a disk eight pack to store data using the redundant array of disks (RAID) data-striping methodology.

nonremovable medium. A recording medium that cannot be added to or removed from a storage device.

nonvolatile storage (NVS). 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*.

organizationally unique identifier (OUI). An IEEE-standards number that identifies an organization with a 24-bit globally unique assigned number referenced by various standards. OUI is used in the family of 802 LAN standards, such as Ethernet and Token Ring.

original equipment manufacturer's information (OEMI). A reference to an IBM guideline for a computer peripheral interface. The interface uses ESA/390 logical protocols over an I/O interface that configures attached units in a multidrop bus topology.

OS/390. The IBM operating system that includes and integrates functions 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 iSeries @server 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.

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.

pSeries. The product name of an IBM @server 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 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*.

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. (1) 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. (2) In ESS documentation, when noted, the term can refer to both S/390 and zSeries storage. See also *zSeries storage*.

SAID. See *system adapter identification number*.

SAM. See *sequential access method*.

SAN. See *storage area network*.

SBCON. See *Single-Byte Command Code Sets Connection*.

screen. The physical surface of a display device upon which information is shown to users.

SCSI. See *small computer system interface*.

SCSI device. A disk drive connected to a host through an I/O interface using the SCSI protocol. A SCSI device is either an initiator or a target. See also *initiator* and *small computer system interface*.

SCSI host systems. Host systems that are attached to the ESS with a SCSI interface. Such host systems run on UNIX, OS/400, Windows NT, Windows 2000, or Novell NetWare operating systems.

SCSI ID. A unique identifier assigned to a SCSI device that is used in protocols on the SCSI interface to identify or select the device. The number of data bits on the SCSI bus determines the number of available SCSI IDs. A wide interface has 16 bits, with 16 possible IDs.

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 TotalStorage Enterprise Storage Server Subsystem Device Driver*.

Seascape architecture. A storage system architecture developed by IBM for open-systems servers and S/390 and zSeries 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 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. (1) A host that provides certain services to other hosts that are referred to as clients. (2) 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.

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-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 multi-processor*.

SNMP. See *simple network management protocol*.

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*.

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. (1) A physical unit that consists of a storage server integrated with one or more storage devices to provide storage capability to a host computer. (2) 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 TotalStorage Enterprise Storage Server Subsystem Device Driver*.

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 multi-processor (SMP). An implementation of a multi-processor 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). (1) Nominally, 1 000 000 000 000 bytes, which is accurate when speaking of bandwidth and disk storage capacity. (2) For ESS cache memory, processor storage, real and virtual storage, a terabyte refers to 2^{40} 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.

TotalStorage. See *IBM TotalStorage*.

TPF. See *transaction processing facility*.

track. A unit of storage on a CKD device that can be formatted to contain a number of data records. See also *home address*, *track-descriptor record*, and *data record*.

track-descriptor record (R0). A special record on a track that follows the home address. The control program uses it to maintain certain information about the track. The record has a count field with a key length of zero, a data length of 8, and a record number of 0. This record is sometimes referred to as R0.

transaction processing facility (TPF). A high-availability, high-performance IBM operating system, designed to support real-time, transaction-driven applications. The specialized architecture of TPF is intended to optimize system efficiency, reliability, and responsiveness for data communication and database processing. TPF provides real-time inquiry and updates to a large, centralized database, where message length is relatively short in both directions, and response time is generally less than three seconds. Formerly known as the Airline Control Program/Transaction Processing Facility (ACP/TPF).

Transmission Control Protocol (TCP). A communications protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

Transmission Control Protocol/Internet Protocol (TCP/IP). (1) A combination of data-transmission protocols that provide end-to-end connections between applications over interconnected networks of different types. (2) A suite of transport and application protocols that run over the Internet Protocol. (GC) See also *Internet Protocol* and *Transmission Control Protocol*.

transparency. See *software transparency*.

TSO. See *time sharing option*.

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.

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). Refer to 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.

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 @server product that emphasizes industry-standard server scalability and self-managing server technologies. It is the successor to the Netfinity family of servers.

Z

z/Architecture. An IBM architecture for mainframe computers and peripherals. The IBM @server zSeries family of servers uses the z/Architecture architecture. It is the successor to the S/390 and 9672 family of servers. See also *Enterprise Systems Architecture/390*.

z/OS. An operating system for the IBM eServer product line that supports 64-bit real storage.

zSeries. (1) An IBM @server family of servers that emphasizes near-zero downtime. (2) IBM enterprise servers based on z/Architecture.

zSeries storage. Storage arrays and logical volumes that are defined in the ESS as connected to zSeries servers. See also *S/390 storage*.

Index

Special characters

/opt/IBMsdd/bin/showvpath command 225

A

about this book xvii
accessing
  AIX
    add a data path volume to a volume group SMIT panel 82
    add a volume group with data path devices SMIT panel 82
    add paths to available data path devices SMIT panel 81
    backup a volume group with data path devices SMIT panel 83
    configure a defined data path device SMIT panel 81
    define and configure all data path devices SMIT panel 81
    display data path device adapter status SMIT panel 80
    display data path device configuration SMIT panel 79
    Display Data Path Device Status SMIT panel 80
    remake a volume group with data path devices SMIT panel 84
    remove a data path device SMIT panel 82
    Remove a Physical Volume from a Volume Group SMIT panel 83
adapter
  firmware level 16, 104
  upgrading firmware level to (sf320A9) 105
adapters 189
  configuring
    Linux 173
    NetWare 189
    Windows 2000 247, 248
    Windows NT 229
    Windows Server 2003 265
  LP70000E 12, 102
adding
  paths
    Windows 2000 host systems 254
    Windows NT 231
    Windows Server 2003 host systems 270
  storage for Windows NT host systems 236
adding paths
  AIX
    to SDD devices volume group 41
addpaths
  utility programs, AIX 85
addpaths command 74, 85
agreement for licensed internal code 311
AIX
  4.3.3
    32–bit applications 17
AIX (continued)
  4.3.3 (continued)
    64–bit applications 17
  5.1.0
    32–bit applications 17
    64–bit applications 17
  accessing
    add a data path volume to a volume group SMIT panel 82
    add a volume group with data path devices SMIT panel 82
    add paths to available data path devices SMIT panel 81
    backup a volume group with data path devices SMIT panel 83
    configure a defined data path device SMIT panel 81
    define and configure all data path devices SMIT panel 81
    display data path device adapter status SMIT panel 80
    display data path device configuration SMIT panel 79
    Display Data Path Device Status SMIT panel 80
    remake a volume group with data path devices SMIT panel 84
    Remove a Physical Volume from a Volume Group SMIT panel 83
    removing a data path device SMIT panel 82
  applications
    32-bit 17
    64-bit 17
  backing-up files belonging to an SDD volume group 78
  changing the path-selection policy 40, 181
  configuring
    volume group for failover protection 70
  configuring SDD 28, 30
  dynamically adding paths to SDD devices of a volume group 41
  error messages 96
    messages for HACMP environment 96
  exporting
    volume group with SDD 76
  extending
    an existing SDD volume group 77
  importing
    volume group with SDD 75
  installing SDD 21
  installing SDDPCM 105
  migrating
    an existing non-SDD volume group to SDD vpath devices in concurrent mode 93
    non-SDD volume group to a SAN Volume Controller for Cisco MDS 9000 SDD multipath volume group in concurrent mode 91

© Copyright IBM Corp. 1999, 2004

AIX *(continued)*
migrating *(continued)*
non-SDD volume group to a SAN Volume Controller SDD multipath volume group in concurrent mode 91
non-SDD volume group to an ESS SDD multipath volume group in concurrent mode 91
MPIO-capable devices 112
recovering
from mixed volume groups 77
removing SDD from a host system 39
removing SDDPCM from a host system 116
restoring files belonging to an SDD volume group 78
SDD utility programs 85
SDD-specific SMIT panels 78
the loss of a device path 72
unconfiguring SDD 31
unsupported environments 10, 101
upgrading 33
upgrading manually 33
verifying SDD 38
verifying SDD installation 22
AIX 5.1.0
32–bit 18
64–bit 18
AIX 5.2.0
32–bit 18
64–bit 18
AIX fibre-channel requirements 14, 103
AIX host system
disk driver 1
protocol stack 1
AIX Hot Plug support 43
AIX trace 95, 137
article
Microsoft Knowledge Base Article Number Q293778
information about removing multipath access to your shared volume 238

B

backing-up AIX files belonging to an SDD volume group 78
BIOS, disabling 229, 248, 265
block disk device interfaces (SDD) 142, 203
boot -r command 223
bootinfo -K command 18

C

cat /proc/modules command 176
cat /proc/scsi/scsi command 177
cat /proc/scsi/xxx/N command 177
cat /proc/sdd command 177
cd /media command 174
cd /mnt command 174
cd /opt/IBMsdd command 175
cd /opt/IBMsdd/bin command 175
cfallvpath 19
cfgmgr
run *n* times where *n* represents the number of paths per SDD device. 74
run for each installed SCSI or fibre-channel adapter 74
cfgmgr command 15, 42, 74, 108
cfgvpath command 179
changing
path-selection policy for AIX 40, 181
path-selection policy for HP 152
path-selection policy for Solaris 212
SDD hardware configuration
HP-UX host systems 151
Solaris hosts 210
to the /dev directory
HP-UX host systems 161
chdev command 73
chgrp command 169, 222
chkconfig - -level X sdd on command 180
chkconfig - -level *X* sdd off command 180
chkconfig - -list sdd command 180
chkvpenv command 176
chmod command 169, 222
chown command 169
command
addpaths 74, 85
bootinfo -K 18
cfgmgr 15, 74, 108
running *n* times for n-path configurations 42, 74
running for each relevant SCSI or FCP adapter 42
chdev 73
datapath disable ports 284
datapath enable ports 285
datapath open device path 286
datapath query adapter 288
datapath query adaptstats 290
datapath query device 70, 74, 291
datapath query devstats 293
datapath query essmap 295
datapath query portmap 296
datapath query wwpn 297
datapath remove adapter 298
datapath remove device 299
datapath remove device path 299
datapath set adapter 301
datapath set device 0 path 0 offline 304
datapath set device N policy rr/fo/lb/df 41, 152, 182, 212
datapath set device path 304
datapath set device policy 303
datapath set qdepth 305
dpovgfix 69, 85
dpovgfix vg-name 42, 73
extendvg 78
extendvg4vp 78, 87
hd2vp and vp2hd 85
hd2vp vg_name 35
HP-UX host system
hd2vp 151
vgexport 158

command *(continued)*
HP-UX host system *(continued)*
vgimport 159
vp2hd 151
installp 15
instfix -i | grep IY10201 15
instfix -i | grep IY10994 15
instfix -i | grep IY11245 15
instfix -i | grep IY13736 15
instfix -i | grep IYl7902 15
instfix -i | grep IYl8070 15
ls -al /unix 18
lscfg -vl fcsN 16, 104
lsdev -Cc disk 16, 109
lsdev -Cc disk | grep 2105 28
lsdev -Cc disk | grep SAN Volume Controller 28
lslpp -l ibmSdd_432.rte 22
lslpp -l ibmSdd_433.rte 22, 23
lslpp -l ibmSdd_510.rte 22, 23
lslpp -l ibmSdd_510nchacmp.rte 22, 23
lspv 34, 42, 71
lsvg -p vg-name 71
lsvgfs 34
lsvpcfg 35, 43, 69, 72, 86
mkdev -l vpathN 31
mksysb restore command 72
mkvg 71
mkvg4vp 71, 87
odmget -q ″name = ioaccess″ CuAt 58
pcmpath disable ports 119
pcmpath enable ports 121
pcmpath open device path 123
pcmpath query adapter 125
pcmpath query adaptstats 126
pcmpath query device 127
pcmpath query devstats 129
pcmpath query essmap 131
pcmpath query portmap 132
pcmpath query wwpn 133
pcmpath set adapter 134
pcmpath set device 0 path 0 offline 135
pcmpath set device path 135
querysn 18, 87
restvg 78
restvg4vp 78
rmdev 74, 75
rmdev -dl dpo -R 34, 39, 116
rmdev -dl fcs*N* -R 16, 109
rmdev -l dpo -R 31
savevg 78
savevg4vp 78
shutdown -rF 15, 109
smitty 34
smitty deinstall 15
smitty device 34
table of, in installation package 19
umount 34
using 117, 283
varyoffvg 28, 34
varyonvg vg_name 35
commands
/opt/IBMsdd/bin/showvpath 225
boot -r 223
cat /proc/modules 176
cat /proc/scsi/scsi 177
cat /proc/scsi/xxx/N 177
cat /proc/sdd 177
cd /media 174
cd /mnt 174
cd /opt/IBMsdd 175
cd /opt/IBMsdd/bin 175
cfgvpath 179
chgrp 169, 222
chkconfig - -level X sdd off 180
chkconfig - -level *X* sdd on 180
chkconfig - -list sdd 180
chkvpenv 176
chmod 169, 222
chown 169
datapath query
adapter 242, 261, 277
device 180, 230, 249, 251, 261, 267, 277
datapath query device 146
datapath set adapter # offline 259, 275
datapath set adapter offline 259, 275
esvpcfg 186
insmod ./vpath.o 176
ls -l 175
lsvpcfg 180
lsvpsd 178
metadb -a <device> 224
metadb -d -f <device> 224
metadb -i 223
metainit 223
metainit d <metadevice number> -t <″vpathNs″ - master device> <″vpathNs″ - logging device> 225
metastat 223, 225
newfs 224
orainst /m 167, 220
pkgrm IBMsdd 224
rmvpath xxx 179
rpm -e IBMsdd command 183
rpm -qi IBMsdd 175, 183
rpm -ql IBMsdd 175, 183
sdd restart 180
sdd start 180
sdd stop 180
showvpath 161, 220, 222, 223, 224
shutdown -i6 -y -g0 223
umount 225
umount /cdrom 210
unmod ./sdd-mod.o 183
commands datapath set adapter # offline 241
commands datapath set adapter offline 241
comments, how to send xxxii
concurrent download of licensed internal code
AIX 47
SDD 7
configuring
additional paths on a Windows NT host system 232

configuring *(continued)*
AIX
cabling storage side switch ports 16, 109
ESS 13, 102
fibre-channel attached devices 14, 103
fibre-channel-attached devices 15, 108
SAN Volume Controller 13
SAN Volume Controller for Cisco MDS 9000 13
volume group for failover protection 70
clusters with SDD
Windows 2000 host system 259
Windows NT host system 241
Windows Server 2003 host system 275
ESS
HP-UX host systems 142
Linux host systems 172
NetWare host systems 189
Solaris host systems 202
Windows 2000 247
Windows NT 229
Windows Server 2003 265
fibre-channel adapters
Linux host systems 173
NetWare host systems 189
Windows 2000 host system 247
Windows NT host system 229
Windows Server 2003 host system 265
SAN File System
Linux host systems 173
SAN Volume Controller
Windows Server 2003 host system 265
SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000
Linux host systems 172
Solaris host systems 202
SCSI adapters
Windows 2000 host systems 248
Windows NT 229
Windows Server 2003 host systems 265
SDD
at system startup 180
Linux host systems 176
NetWare host systems 191
Solaris host systems 210
Windows NT host system 230
SDD for AIX host 28
SDD on AIX 30
configuring a vpath device to the Available state 75
configuring all vpath devices to the Available state 75
conversion script
vp2hd 31
conversion scripts
hd2vp 85
vp2hd 85
creating
device node for the logical volume device in an HP-UX host systems 161
directory in /dev for the volume group in an HP-UX host systems 161
file system on the volume group in an HP-UX host systems 162
creating *(continued)*
logical volume in an HP-UX host systems 161
new logical volumes in an HP-UX host systems 160
physical volume in an HP-UX host systems 161
volume group in an HP-UX host systems 161
customizing
Network File System file server 165
Oracle 166, 219
standard UNIX applications 160, 217

D

database managers (DBMS) 203
datapath
disable porst command 284
enable ports command 285
open device path command 286
query
adapter command 242, 261, 277
device command 180, 230, 242, 261, 277
query adapter command 288
query adaptstats command 290
query device command 291
query devstats command 293
query essmap command 295
query portmap command 296
query set adapter command 301
query wwpn command 297
remove adapter command 298
remove device path command 299
set adapter # offline command 241, 259, 275
set adapter offline command 241, 259, 275
set device path command 304
set device policy command 303
set qdepth command 305
datapath query adapter 153, 298, 299
datapath query device command 70, 74
datapath remove adapter 298
datapath remove device 299
datapath set device 0 path 0 offline command 304
datapath set device N policy rr/fo/lb/df command 41, 152, 182, 212
datapath set qdepth 305
definitions 313
determining
AIX
adapter firmware level 16, 104
major number of the logical volume device for an HP-UX host systems 160
size of the logical volume for an HP-UX host systems 163
device driver 202
devices.fcp.disk.ibm2062.rte 13
devices.fcp.disk.ibm2105.rte 9
devices.fcp.disk.ibm2145.rte 9, 13
devices.scsi.disk.ibm2105.rte 9
displaying
AIX
ESS vpath device configuration 69
current version of SDD
Windows 2000 250

displaying *(continued)*
current version of SDD *(continued)*
Windows Server 2003 268
current version of the SDD
Windows NT 239
dpovgfix command 69, 85
dpovgfix vg-name command 42, 73
dynamic I/O load-balancing 6
Dynamically opening an invalid or close_dead path 43
Dynamically removing paths 46
Dynamically replacing adapters
AIX Hot Plug support 43
different type replacement 44
same type replacement 43

E

enhanced data availability 3
error log messages
AIX
VPATH_DEVICE_ONLINE 96
error messages
AIX
messages for HACMP environment 96
VPATH_DEVICE_OFFLINE 96
VPATH_PATH_OPEN 96
VPATH_XBUF_NOMEM 96
error messages for ibmSdd_433.rte installation package for SDD
AIX
VPATH_FAIL_RELPRESERVE 96
VPATH_OUT_SERVICE 96
VPATH_RESV_CFLICT 96
error recovery
HP-UX 11.0 host system 156
HP-UX 11i host system 156
Solaris host system 214
error recovery and retry policy
algorithms
Windows 2000 host systems 257
Windows NT host systems 239
Windows Server 2003 host systems 273
error recovery and retry policy, for Windows 2000 host systems 257
error recovery and retry policy, for Windows NT host systems 239
error recovery and retry policy, for Windows Server 2003 host systems 273
error recovery policy
algorithms
Linux host systems 184
error recovery policy, for Linux host systems 184
ESS
AIX
displaying vpath device configuration 69
configuring for HP 142
configuring for Linux 172
configuring for NetWare 189
configuring for Solaris 202
configuring on
Windows 2000 247
ESS *(continued)*
configuring on *(continued)*
Windows Server 2003 265
configuring on Windows NT 229
ESS devices (hdisks) 90
ESS LUNs 90
exporting a volume group with SDD, AIX 76
extending an existing SDD volume group, AIX 77
extendvg command 78
extendvg4vp command 78, 87

F

failover 6
failover protection, AIX
creating a volume group from a single-path vpath device 72
losing 71
manually deleted devices and running the configuration manager 74
providing load-balancing and failover protection 69
side effect of running the disk change method 72
the loss of a device path 72
when it doesn't exist 69
fibre-channel adapters
configuring
Linux host systems 173
NetWare host systems 189
configuring for Windows 2000 247
configuring for Windows Server 2003 265
supported
HP-UX host systems 141
Linux host systems 171
Solaris host systems 201
Windows 2000 host systems 246
Windows NT host systems 228
Windows Server 2003 host systems 264
supported on AIX host systems 12, 102
supported on NetWare host systems 189
fibre-channel device drivers
configuring for AIX 14, 103
devices.common.IBM.fc 14, 103
devices.fcp.disk 14, 103
devices.pci.df1000f7 14, 103
installing for AIX 14, 103
supported on AIX host systems 12, 102
supported on NetWare host systems 189

G

glossary 313

H

HACMP
concurrent mode 53
hd2vp conversion script 59
node failover 67
nonconcurrent mode 53
persistent reserve 58
recovering paths 67

HACMP *(continued)*
SDD persistent reserve attributes 58
sddsrv 57
software support for concurrent mode 56
software support for nonconcurrent mode 54
special requirements 58
hardware configuration
changing
HP-UX host systems 151
Solaris host systems 210
hardware requirements
HP
host systems 141
Linux host systems 171
Solaris host systems 201
hd2vp and vp2hd command 85
hd2vp command
HP-UX host system 151
hd2vp vg_name command 35
hdisk device
chdev 72
modify attributes 72
healthcheck 113
High Availability Cluster Multi-Processing (HACMP) 53
HP
changing the path-selection policy 152
SCSI disk driver (sdisk) 142
unsupported environments 141
HP host system
disk driver 1
protocol stack 1
HP-UX
disk device drivers 146, 160
LJFS file system 165
operating system 141
HP-UX 11.0
32-bit 142, 143
64-bit 142, 143
HP-UX 11i
32-bit 142, 143
64-bit 142, 143
HP-UX host systems
changing
SDD hardware configuration 151
to the /dev directory 161
creating
a file system on the volume group 162
a logical volume 161
a volume group 161
device node for the logical volume device 161
directory in /dev for the volume group 161
new logical volumes 160
physical volume 161
determining
major number of the logical volume 160
size of the logical volume 163
installing Oracle 166
installing SDD 144
converting an Oracle installation from sdisk 168
on a Network File System file server 165
HP-UX host systems *(continued)*
installing SDD *(continued)*
on a system that already has Network File System file server 166
on a system that already has Oracle 168
using a file system 168
using raw partitions 168
mounting the logical volume 162
re-creating
existing logical volume 163
logical volume 164
physical volume 161, 163
volume group 163
removing
existing logical volume 162
existing volume group 163
logical volumes 162
SDD 141
setting the correct timeout value for the logical volume manager 164
setting up Network File System for the first time 165
setting up Oracle
using a file system 166
using raw partitions 166
standard UNIX applications 160
understanding how SDD works 142
upgrading SDD 143, 149
using applications with SDD 160

I

ibm2105.rte 13
ibm2105.rte ESS package 11, 188
ibmSdd_433.rte installation package
for SDD 1.2.2.0
removing 58
for SDD 1.3.2.0. vpath devices
unconfiguring 58
importing a volume group with SDD, AIX 75
insmod ./sdd-mod.o command 176
installation package
AIX 21
devices.fcp.disk.ibm2105.mpio.rte 105
devices.sdd.*nn*.rte 21
devices.sdd.43.rte 17, 31
devices.sdd.51.rte 17, 31
devices.sdd.52.rte 17
devices.sddpcm.52b.rte 116
ibmSdd_432.rte 22, 40, 54, 91
ibmSdd_433.rte 22, 23, 40, 54, 58, 76, 91, 96
ibmSdd_510.rte 23, 40, 54
ibmSdd_510nchacmp.rte 40, 54
installing
additional paths on a Windows NT host system 232
AIX
fibre-channel device drivers 14, 103
planning 9, 99
SDD 21
SDDPCM 105

installing *(continued)*
an Oracle installation from sdisk on an HP-UX host systems 168
converting an Oracle installation from sdisk on a Solaris host system 221
NetWare
planning 187
Oracle
HP-UX host systems 166
Solaris host systems 219
SDD
HP-UX host systems 141, 144
Linux host systems 171, 174
NetWare host systems 191
Solaris host systems 201, 204
Windows 2000 host system 245
Windows NT host system 227, 230
Windows Server 2003 host system 263
SDD 1.4.0.0 (or later)
Windows 2000 host system 248
SDD 1.5.0.0 (or later)
Windows Server 2003 host system 266
SDD on a Network File System file server on a Solaris host system 217
SDD on a Network File System file server on an HP-UX host systems 165
SDD on a system that already has Network File System file server 217
SDD on a system that already has Network File System file server on an HP-UX host systems 166
SDD on a system that already has Oracle on a Solaris host system 220
SDD on a system that already has Oracle on an HP-UX host systems 168
SDD on a system that already has Solstice DiskSuite in place on a Solaris host system 223
Solaris Volume Manager for the first time on a Solaris host system 223
using a file system on a Solaris host system 220
using a file system on an HP-UX host systems 168
using raw partitions on a Solaris host system 221
using raw partitions on an HP-UX host systems 168
vpath on a system that already has UFS logging in place on a Solaris host system 225
installp command 15
instfix -i | grep IY10201 command 15
instfix -i | grep IY10994 command 15
instfix -i | grep IY11245 command 15
instfix -i | grep IY13736 command 15
instfix -i | grep IYl7902 command 15
instfix -i | grep IYl8070 command 15

K

KB 130, 294

L

licensed internal code
agreement 311
Linux
unsupported environments 172
Linux host system
disk driver 1
protocol stack 1
Linux host systems
configuring
ESS 172
fibre-channel adapters 173
SAN File System 173
SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000 172
SDD 176, 180
error recovery policy 184
installing SDD 171, 174
maintaining vpath device configuration persistency 181
partitioning SDD vpath devices 185
preparing SDD installation 172
removing SDD 183
upgrading SDD 174
using
multiple command for SDD configuration 176
single command for SDD configuration 180
standard UNIX applications 186
verifying SDD installation 175
load-balancing, AIX 69
loading SDD on Linux 176
loading SDD on NetWare 191
loading SDD on Solaris 210
logical volume manager 203
losing failover protection, AIX 71
ls -al /unix command 18
ls -l command 175
lscfg -vl fcsN command 16, 104
lsdev -Cc disk | grep 2105 command 28
lsdev -Cc disk | grep SAN Volume Controller command 28
lsdev -Cc disk command 16, 109
lslpp -l '*Sdd*' command 22
lslpp -l ibmSdd_432.rte command 22
lslpp -l ibmSdd_433.rte command 22, 23
lslpp -l ibmSdd_510.rte command 22, 23
lslpp -l ibmSdd_510nchacmp.rte command 22, 23
lspv command 34, 42, 71
lsvg -p vg-name command 71
lsvgfs command 34
lsvpcfg 19
lsvpcfg command 35, 43, 69, 72, 86, 180, 186
lsvpcfg utility programs, AIX 86
lsvpd command 178

M

maintaining vpath device configuration persistency, for Linux host systems 181
Manual Exclusion of ESS devices from the SDD configuration 18
metadb -a <device> command 224
metadb -d -f <device> command 224
metadb -i command 223

metainit command 223
metainit d <metadevice number> -t <″vpathNs″ - master device> <″vpathNs″ - logging device> command 225
metastat command 223, 225
migrating
  AIX
    an existing non-SDD volume group to SDD vpath devices in concurrent mode 93
    non-SDD volume group to a SAN Volume Controller SDD multipath volume group in concurrent mode 91
    non-SDD volume group to an ESS SDD multipath volume group in concurrent mode 91
    non-SDD volume group to aSAN Volume Controller for Cisco MDS 9000 SDD multipath volume group in concurrent mode 91
mirroring logical volumes 93
mkdev -l vpath*N* command 31
mksysb restore command 72
mkvg command 71
mkvg4vp command 71, 87
modifying multipath storage configuration to the ESS, Windows NT host system 236
mounting the logical volume, HP 162

N

NetWare
  example command output 195
NetWare host systems
  configuring
    ESS 189
    fibre-channel adapters 189
    SDD 191
  error logging 194
  error reporting 194
  installing SDD 191
  preparing SDD installation 189
  removing SDD 195
  supported environments 188
  unsupported environments 188
newfs command 224
notices
  licensed internal code 311
  notices statement 309

O

odmget -q ″name = ioaccess″ CuAt command 58
orainst /m command 167, 220

P

partitioning SDD vpath devices, for Linux host systems 185
path-failover protection system 6
path-selection policy
  changing 41, 152, 182, 212
  default (optimized) 192
  failover only 41, 152, 181, 193, 212
  load balancing 152, 181, 192, 212
path-selection policy *(continued)*
  load-balancing 41
  round robin 41, 152, 182, 192, 212
pcmpath
  disable ports command 119
  enable ports command 121
  open device path command 123
  query adapter command 125
  query adaptstats command 126
  query device command 127
  query devstats command 129
  query essmap command 131
  query portmap command 132
  query set adapter command 134
  query wwpn command 133
  set device path command 135
pcmpath set device 0 path 0 offline command 135
Persistent Reserve command set 58
pkgrm IBMsdd command 224
planning
  AIX
    adapter firmware level 16, 104, 105
    ESS 13, 102
    fibre-channel attached devices 14, 103
    fibre-channel device drivers 14, 103
    fibre-channel-attached devices 15, 108
    preparing 12, 102
    SAN Volume Controller 13
    SAN Volume Controller for Cisco MDS 9000 13
  AIX installation 9, 99
  ESS
    HP-UX host systems 142
    Linux host systems 172
    NetWare host systems 189
    Solaris host systems 202
    Windows 2000 host system 247
    Windows NT host system 229
    Windows Server 2003 host system 265
  fibre-channel adapters
    Windows 2000 host system 247
    Windows NT host system 229
    Windows Server 2003 host system 265
  hardware and software requirements on a Solaris host system 201
  hardware and software requirements on an HP-UX host systems 141
  hardware requirements
    ESS 101
    Fibre adapters and cables 101
  hardware requirements, AIX
    ESS 9
    Fibre channel adapters and cables 9
    Host system 9
    SAN Volume Controller 9
    SCSI adapters and cables 9
  hardware requirements, SDDPCM
    ESS 101
    Fibre adapters and cables 101
    Host system 101
  hardware requirements, Windows 2000
    ESS 245

planning *(continued)*
hardware requirements, Windows NT
ESS 227
hardware requirements, Windows Server 2003
ESS 263
host system requirements, AIX 10
ESS 11, 101
Fibre 11
SAN Volume Controller 11
SAN Volume Controller for Cisco MDS 9000 11
SCSI 11
host system requirements, NetWare 187
ESS 188
Fibre 188
SCSI 188
host system requirements, SDDPCM 101
Fibre 101
host system requirements, Windows 2000
ESS 245
host system requirements, Windows NT
ESS 228
host system requirements, Windows Server 2003
ESS 263
installation of SDD
HP-UX host systems 142
Solaris host systems 203
NetWare installation 187
preparing
Solaris host systems 202
preparing for SDD installation on an HP-UX host systems 142
SAN File System
Linux host systems 173
SAN Volume Controller
Windows Server 2003 host system 265
SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000
Linux host systems 172
Solaris host systems 202
SCSI adapters
Windows NT host systems 229
SDD
HP-UX host systems 141
Linux host systems 171, 172
NetWare host systems 189
Solaris host systems 201
Windows 2000 host system 247
Windows NT host system 227
Windows Server 2003 host system 264
software requirements
Windows 2000 operating system 245
Windows NT operating system 227
Windows Server 2003 operating system 263
software requirements, AIX
AIX operating system 9
ibm2105.rte ESS package 9
SCSI and fibre-channel device drivers 9
software requirements, AIX 5.2.0 ML3
fibre-channel device drivers 101
software requirements, SDDPCM
AIX 5.2.0 ML3 operating system 101
planning *(continued)*
Windows 2000
ESS 247
Windows Server 2003
ESS 265
Windows Server 2003 host system
SAN Volume Controller 265
postinstallation of SDD
HP-UX host systems 146
Solaris host systems 208
preparing
AIX
SDD installation 12
SDDPCM installation 102
configure on AIX 28
SDD
HP-UX host systems 142
Linux host systems 172
NetWare host systems 189
Windows 2000 installation 247
Windows NT host system 229
Windows Server 2003 installation 264
SDD installation
Solaris host systems 202
providing
AIX
failover protection 69
load-balancing 69
pSeries 690 with static LPARs configured 52
pvid 92
PVID 71

Q

qdepth_enable 29

R

raw
device interface (sd) 203
device interface (sdisk) 142
re-creating
existing logical volume
on a HP-UX host systems 163
physical volume
on an HP-UX host systems 161
the logical volume
on an HP-UX host systems 164
the physical volume
on a HP-UX host systems 163
the volume group
on an HP-UX host systems 163
recovering from mixed volume groups 77
recovering from mixed volume groups, AIX 77
recovery procedures for HP 162, 165
removing
existing logical volume
on an HP-UX host system 162
existing volume group
on an HP-UX host systems 163

removing *(continued)*
logical volumes
on an HP-UX host systems 162
SDD
Linux host systems 183
NetWare host systems 195
Windows 2000 host system 252
Windows NT host system 238
Windows Server 2003 host system 269
SDD from an AIX host 39
SDD from an AIX host system 39
SDDPCM from an AIX host 116
SDDPCM from an AIX host system 116
removing SDD in a two-node cluster environment 253, 269
Replacing manually excluded devices in the SDD configuration 19
requirements
ESS
Windows 2000 host system 245
Windows NT 228
Windows Server 2003 host system 263
hardware
ESS 101
Fibre adapters and cables 101
hardware and software
Linux host systems 171
hardware and software on a Solaris host system 201
hardware and software, HP 141
hardware, AIX
ESS 9
Fibre channel adapters and cables 9
Host system 9
SAN Volume Controller 9
SCSI adapters and cables 9
hardware, SDDPCM
ESS 101
Fibre adapters and cables 101
Host system 101
hardware, Windows 2000
ESS 245
hardware, Windows NT
ESS 227
hardware, Windows Server 2003
ESS 263
host system, AIX 10
ESS 11, 101
Fibre 11
SAN Volume Controller 11
SAN Volume Controller for Cisco MDS 9000 11
SCSI 11
host system, NetWare 187
ESS 188
Fibre 188
SCSI 188
host system, SDDPCM 101
Fibre 101
software
Windows 2000 operating system 245
Windows NT operating system 227
requirements *(continued)*
software *(continued)*
Windows Server 2003 operating system 263
software, AIX
AIX operating system 9
ibm2105.rte ESS package 9
SCSI and fibre-channel device drivers 9
software, SDDPCM
AIX 5.2.0 ML3 operating system 101
devices.sddpcm.52b.rte ESS package 101
restoring
AIX
files belonging to an SDD volume group 78
restvg command 78
restvg4vp command 78
reviewing the existing SDD configuration information, Windows NT 231, 236
rmdev -dl dpo -R command 34, 39, 116
rmdev -dl fcsN -R command 16, 109
rmdev -l dpo -R command 31
rmdev command 74, 75
rmvpath xxx command 179
rpm -e IBMsdd command 183
rpm -qi IBMsdd command 175, 183
rpm -ql IBMsdd command 175, 183

S

SAN File System
configuring for Linux 173
SAN Volume Controller
configuring on
Windows Server 2003 host system 265
SAN Volume Controller for Cisco MDS 9000 Preferred Node path selection algorithm 40, 152, 213, 234, 257
SAN Volume Controller Preferred Node path selection algorithm 40, 152, 213, 257
SAN Volume Controller and SAN Volume Controller for Cisco MDS 9000
configuring for Linux 172
configuring for Solaris 202
savevg command 78
savevg4vp command 78
SCSI adapter support
AIX host systems 11
HP-UX host systems 141
NetWare host systems 188
Solaris host systems 201
Windows 2000 host system 246
Windows NT host system 228
Windows Server 2003 host system 264
SCSI-3 Persistent Reserve command set 58
SDD
architecture 1
displaying the current version on Windows 2000 250
displaying the current version on Windows Server 2003 268
how it works on an HP-UX host system 142
how it works on Solaris 202
installation scenarios 143

SDD *(continued)*
installing
AIX 9
HP-UX host system 141
Linux 171
NetWare 187
Solaris host systems 201
Windows 2000 host system 245, 248
Windows NT 227
Windows Server 2003 host system 263, 266
introduction 1
overview 1
postinstallation of SDD
HP-UX host systems 146
postinstallation on Solaris host systems 208
removing SDD on Windows NT 238
uninstalling
HP-UX host systems 150
upgrading
HP-UX host systems 143
Windows 2000 250
Windows Server 2003 268
using applications
with SDD on HP Network File System file server 165
with SDD on HP-UX standard UNIX applications 160
with SDD on HP, Oracle 166
with SDD on Linux standard UNIX applications 186
with SDD on Solaris Network File System file Server 217
with SDD on Solaris standard UNIX applications 217
with SDD on Solaris, Oracle 219
verifying additional paths to SDD devices 232, 255, 272
verifying configuration 38
Web site xviii
SDD configuration
checking 38
SDD devices
reconfiguring 43
sdd restart command 180
SDD server daemon 279
AIX host system 20, 48, 183
HP-UX host system 155
Solaris host system 216
Windows 2000 host system 253
Windows NT host system 238
Windows Server 2003 host system 270
sdd start command 180
sdd stop command 180
SDD utility programs, AIX 85
SDD vpath devices 90
sddpcm 279
SDDPCM
installing
AIX 99
updating package 109
SDDPCM server 109
SDDPCM server daemon 279
sddserver.rte
AIX host system 16
sddsrv 57, 279
AIX host system 20, 48, 183
HP-UX host system 155
port binding 279
Solaris host system 216
trace 280
Windows 2000 host system 253
Windows NT host system 238
Windows Server 2003 host system 270
sddsrv for ESS Expert
AIX host system 16
setting up
correct timeout value for the logical volume manager on an HP-UX host system 164
Network File System for the first time on a Solaris host system 217
Network File System for the first time on an HP-UX host system 165
Oracle using a file system
HP-UX host system 166
Solaris host system 219
Oracle using raw partitions
HP-UX host system 166
Solaris host system 219
UFS logging on a new system on a Solaris host system 224
showvpath command 161, 220, 222, 223, 224
shutdown -i6 -y -g0 command 223
shutdown -rF command 15, 109
smit
definition 21, 107
SMIT
configuring
SDD for Windows NT host system 230
smitty
definition 21, 107
smitty command 34
smitty deinstall command 15
smitty device command 34
software requirements
for SDD on HP 141
for SDD on Linux 171
for SDD on Solaris 201
Solaris
changing the path-selection policy 212
host system
upgrading the Subsystem Device Driver on 203
operating system
upgrading SDD 201
sd devices 216
supported environments 201
UFS file system 217
unsupported environments 202
Solaris disk device drivers 203
Solaris host systems
changing SDD hardware configuration 210
configuring
SDD 210

Solaris host systems *(continued)*
installing
Solaris Volume Manager for the first time 223
vpath on a system that already has UFS logging in place 225
installing Oracle 219
installing SDD 204
converting an Oracle installation from sdisk 221
Network File System file server 217
system that already has Network File System file server 217
system that already has Oracle 220
system that already has Solstice DiskSuite in place 223
using a file system 220
using raw partitions 221
Oracle 219
SDD 201
SDD postinstallation 208
setting up
Network File System for the first time 217
UFS logging on a new system 224
Solaris Volume Manager 222
Solstice DiskSuite 222
standard UNIX applications 217
understanding how SDD works 202
upgrading SDD 210
using applications with SDD 216, 218
Solaris SCSI disk driver 202
Subsystem device driver, see SDD. 222
Sun host system
disk driver 1
protocol stack 1
support for Windows 2000 259
support for Windows NT 241
support for Windows Server 2003 274
supported environments
NetWare host systems 188
Solaris 201
synchronizing logical volumes 93
System Management Interface Tool (SMIT) 21, 107
definition 21, 107
using for configuring 28
using to access the Add a Data Path Volume to a Volume Group panel on AIX host 82
using to access the Add a Volume Group with Data Path Devices panel on AIX host 82
using to access the backup a Volume Group with Data Path Devices on AIX host 83
using to access the Configure a Defined Data Path Device panel on AIX host 81
using to access the Define and Configure All Data Path Devices panel on AIX host 81
using to access the Display Data Path Device Adapter Status panel on AIX host 80
using to access the Display Data Path Device Configuration panel on AIX host 79
using to access the Display Data Path Device Status panel on AIX host 80
using to access the Remake a Volume Group with Data Path Devices on AIX host 84
System Management Interface Tool (SMIT) *(continued)*
using to access the Remove a Data Path Device panel on AIX host 82
using to access the Remove a Physical Volume from a Volume Group panel on AIX host 83
using to back up a volume group with SDD on AIX host 78
using to backup a volume group with SDD on AIX host 83
using to create a volume group with SDD on AIX host 70
using to display the ESS vpath device configuration on AIX host 69
using to display the SAN Volume Controller for Cisco MDS 9000 vpath device configuration on AIX host 69
using to display the SAN Volume Controller vpath device configuration on AIX host 69
using to export a volume group with SDD on AIX host 76
using to extend an existing SDD volume group on AIX host 77
using to import a volume group with SDD on AIX host 75
using to remove SDD from AIX host 39
using to remove SDDPCM from AIX host 116
using to restore a volume group with SDD on AIX host 78, 84
using to unconfigure SDD devices on AIX host 31
using to verify SDD configuration on AIX host 38

T

trademarks 310

U

umount
/cdrom command 210
command 225
umount command 34
unconfiguring a SDD device to Defined state 75
unconfiguring all SDD devices to Defined state 75
unconfiguring SDD on AIX 31
understanding
how SDD works for HP-UX host systems 142
how SDD works for Solaris host systems 202
unmod ./sdd-mod.o command 183
unsupported environments
AIX 10, 101
HP 141, 142
Linux 172
NetWare host systems 188
Solaris 202
Windows 2000 245
Windows NT 227
Windows Server 2003 263
Updating SDD using a PTF 35
upgrading
AIX
adapter firmware level 105

upgrading *(continued)*
AIX *(continued)*
manually 33
SDD
HP-UX host system 149
Linux host system 174
Solaris host system 210
Windows 2000 host system 250
Windows NT host system 235
Windows Server 2003 host system 268
upgrading automatically
SDD 32
upgrading manually
SDD
for AIX 4.3.2 33
for AIX 4.3.3 33
for AIX 5.1.0 33
for AIX 5.2.0 33
upgrading to SDD 1.3.3.3 (or later) in a two-node cluster environment 251
using
HP-UX applications with SDD 160
Linux
single command for SDD configuration 180
standard UNIX applications 186
Solaris applications with SDD 216
using command 117, 283
using ESS devices directly, AIX 90
using ESS devices through AIX LVM 91
using SAN Volume Controller devices through AIX LVM 91
using SAN Volume Controller for Cisco MDS 9000 devices through AIX LVM 91
using the datapath commands 283
using the pcmpath commands 117
using the SDDPCM trace function, AIX 137
using the trace function, AIX 95
utility programs, AIX
addpaths 85
dpovgfix 85
extendvg4vp 87
hd2vp and vp2hd 85
lsvpcfg 86
mkvg4vp 87
using ESS devices directly 90
using ESS devices through AIX LVM 91
using SAN Volume Controller devices through AIX LVM 91
using SAN Volume Controller for Cisco MDS 9000 devices through AIX LVM 91
using the SDDPCM trace function 137
using the trace function 95
utility programs, HP
hd2vp 151
vp2hd 151

V

varyoffvg command 28, 34
varyonvg vg_name command 35
verifying
additional paths are installed correctly
Windows 2000 host system 255
Windows NT host system 232
Windows Server 2003 host system 272
AIX
configuring SDD 38
SDD installation 22
SDD installation
Linux host systems 175
Veritas Volume Manager
Command Line Interface for Solaris Web site 218
System Administrator's Guide Web site 218
vgexport command
HP-UX host system 158
vgimport command
HP-UX host system 159
volume group
mixed
how to fix problem 73
mixed volume groups
dpovgfix vg-name 73
volume groups on AIX 71
vp2hd command
HP-UX host system 151

W

Web site
AIX APARs, maintenance level fixes and microcode updates 11, 101
information about removing multipath access to your shared volume
Multiple-Path Software May Cause Disk Signature to Change (Knowledge Base Article Number Q293778) 238
information on the fibre-channel adapters that can be used on your AIX host 12, 102
information on the fibre-channel adapters that can be used on your NetWare host 189
information on the SCSI adapters that can attach to your AIX host 11
information on the SCSI adapters that can attach to your NetWare host 188
NetWare APARs, maintenance level fixes and microcode updates 187
SDD xviii
Web sites
HP-UX documentation 162, 165
information about
SCSI adapters that can attach to your Windows 2000 host system 246
SCSI adapters that can attach to your Windows NT host system 228
SCSI adapters that can attach to your Windows Server 2003 host system 264
Windows 2000
unsupported environments 245
Windows 2000 host system
adding
paths to SDD devices 254

Windows 2000 host system *(continued)*
clustering special considerations 259
configuring
cluster with SDD 259
ESS 247
fibre-channel adapters 247
SCSI adapters 248
disk driver 1
displaying the current version of the SDD 250
installing SDD 1.4.0.0 (or later) 248
path reclamation 259
protocol stack 1
removing SDD 252
SDD 245
support for clustering 259
upgrading SDD 250
verifying
additional paths to SDD devices 255
Windows 2000 host systems
error recovery and retry policy 257
Windows NT
adding
paths to SDD devices 231
unsupported environments 227
Windows NT host system
adding
multipath storage configuration to the ESS 236
new storage to existing configuration 236
clustering special considerations 241
configuring
additional paths 232
clusters with SDD 241
SDD 230
disk driver 1
displaying the current version of the SDD 239
installing
additional paths 232
SDD 230
modifying multipath storage configuration to the ESS 236
path reclamation 241
protocol stack 1
removing SDD 238
reviewing existing SDD configuration information 231, 236
SDD 227
support for clustering 241
upgrading
SDD 235
verifying
additional paths to SDD devices 232
new storage is installed correctly 237
Windows NT host systems
error recovery and retry policy 239
Windows Server 2003
unsupported environments 263
Windows Server 2003 host system
adding
paths to SDD devices 270
clustering special considerations 275
Windows Server 2003 host system *(continued)*
configuring
cluster with SDD 275
ESS 265
fibre-channel adapters 265
SAN Volume Controller 265
SCSI adapters 265
displaying the current version of the SDD 268
installing SDD 1.5.0.0 (or later) 266
path reclamation 275
removing SDD 269
SDD 263
support for clustering 274
upgrading SDD 268
verifying
additional paths to SDD devices 272
Windows Server 2003 host systems
error recovery and retry policy 273

Readers' Comments — We'd Like to Hear from You

IBM TotalStorage Enterprise Storage Server
IBM TotalStorage SAN Volume Controller
IBM TotalStorage SAN Volume Controller for Cisco MDS 9000
Subsystem Device Driver User's Guide

Publication No. SC26-7637-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name _____ Address _____

Company or Organization _____

Phone No. _____

Readers' Comments — We'd Like to Hear from You
SC26-7637-00

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

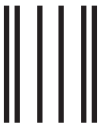

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department GZW
9000 South Rita Road
Tucson, AZ 85744-0001

Fold and Tape | **Please do not staple** | Fold and Tape

Cut or Fold
Along Line

Printed in USA

SC26-7637-00

Spine information:

IBM TotalStorage Enterprise Storage Server
IBM TotalStorage SAN Volume Controller
IBM TotalStorage SAN Volume Controller for Cisco MDS 9000

Subsystem Device Driver User's Guide