



Subsystem Device Driver User's Guide for the IBM TotalStorage Enterprise Storage Server and the IBM TotalStorage SAN Volume Controller

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.

Note

Before using this information and the product it supports, read the information in “Notices” on page 229.

First Edition (June 2003)

This edition replaces SC26-7478-02, and it includes information that specifically applies to the Subsystem Device Driver (SDD) for:

- Version 1 Release 4 Modification 0 Level 0 (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 4 (or later) for HP-UX 11.0 and HP-UX11i
- Version 1 Release 4 Modification 0 Level 0 (or later) for Linux Red Hat 7.x, Red Hat AS2.1, SuSE SLES7 and SuSE SLES8 (Intel i686 only)
- Version 1.0.0.g (or later) for Novell
- Version 1 Release 4 Modification 0 Level 6 (or later) for Solaris 2.6, Solaris 7, Solaris 8
- Version 1 Release 4 Modification 0 Level 0 (or later) for Microsoft[®] Windows NT[®] 4.0 Service Pack 6A or higher
- Version 1 Release 4 Modification 0 Level 0 (or later) for Windows[®] 2000 Service Pack 2 or higher
- Version 1 Release 4 Modification 0 Level 0 (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, 2002, 2003. All rights reserved.**

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

Contents

Figures	xi
Tables	xiii
About this book	xv
Who should use this book	xv
Command syntax conventions	xv
Highlighting conventions	xv
Special characters conventions	xv
Summary of changes	xvi
New information	xvi
Modified information	xvi
Related information	xvii
The ESS library	xvii
The SAN Volume Controller Library	xix
Other IBM publications	xix
Ordering IBM publications	xxvi
Non-IBM publications	xxvi
How to send your comments	xxvii
Chapter 1. Overview of SDD	1
The SDD architecture	1
Enhanced data availability	4
Dynamic I/O load balancing	5
Automatic path-failover protection	5
Concurrent download of ESS licensed internal code	6
Chapter 2. Using SDD on an AIX host system	7
Verifying the hardware and software requirements	7
Hardware	7
Software	7
Unsupported environments	8
Host system requirements	8
Preparing for SDD installation	10
Configuring the ESS	10
Configuring the SAN Volume Controller	11
Installing the AIX Fibre Channel device drivers	11
Uninstalling the AIX Fibre Channel device drivers	12
Configuring Fibre Channel-attached devices	13
Removing Fibre Channel-attached devices	13
Verifying the adapter firmware level	13
Upgrading the adapter firmware level	14
Determining if the sddServer for Expert is installed	14
Planning for SDD installation on a pSeries 690 server LPAR	15
Determining the installation package	15
Determining the installation type	18
Installing SDD	18
Verifying the currently installed version of SDD for SDD 1.3.3.11 (or earlier)	20
Verifying your currently installed version of SDD for SDD 1.4.0.0 (or later)	21
Maximum number of LUNs for a configuration	22
Preparing to configure SDD	24
Configuring SDD	24
Unconfiguring Subsystem Device Driver	25

Migrating or upgrading SDD packages automatically without system restart	26
Preconditions for migration or upgrade	26
Procedures for migration or upgrade	26
Error recovery for migration or upgrade	27
Adding and configuring more ESS LUNs with AIX 5.2.0	27
Migrating or upgrading the SDD package during an AIX OS upgrade	28
Updating SDD packages by applying a program temporary fix	29
Committing or Rejecting a PTF Update	30
Verifying the SDD configuration	30
Changes to the Object Data Manager (ODM) database	31
Displaying the device model	32
Displaying the serial number	32
Removing SDD from an AIX host system	33
SAN Volume Controller Preferred Node path-selection algorithm	34
Dynamically changing the SDD path-selection policy algorithm	34
datapath set device policy command	35
Dynamically adding paths to SDD devices of a volume group	35
Adding paths from AIX 4.3.2 (or later) host systems	35
Dynamically opening an invalid or close_dead path	36
Adding more paths to the existing vpaths	36
Using ESS concurrent download of licensed internal code	37
SDD server daemon	37
Verifying if the SDD server has started	37
Starting the SDD server manually	38
Changing to a different port number for the SDD server	38
Stopping the SDD server	38
PTFs for APARs on AIX with Fibre Channel and the SDD server	39
Understanding SDD 1.3.2.9 (or later) support for single-path configuration for ESS	40
Understanding SDD error recovery policies	40
Understanding SDD support for pSeries 690 with static LPARs configured	42
Understanding SDD support for High Availability Cluster Multi-Processing	42
HACMP and sddsrv	48
SDD persistent reserve attributes	48
Special requirements	49
HACMP RAID concurrent-mode volume groups and enhanced concurrent-capable volume groups	51
HACMP RAID concurrent-mode volume groups	51
Enhanced concurrent-capable volume groups	55
Recovering paths that are lost during HACMP node failover	58
Supporting enhanced concurrent mode in an HACMP environment	58
Managing secondary-system paging space	59
Listing paging spaces	59
Adding a paging space	59
Removing a paging space	59
Providing load-balancing and failover protection	60
Displaying the ESS or SAN Volume Controller vpath device configuration	60
Configuring volume groups for failover protection	61
Importing volume groups with SDD	62
Importing volume groups under HACMP	63
Exporting a volume group with SDD	63
Losing failover protection	64
Recovering from mixed volume groups	68
Extending an existing SDD volume group	68
Backing up all files belonging to an SDD volume group	69
Restoring all files belonging to an SDD volume group	69

SDD-specific SMIT panels	69
SDD utility programs	76
Persistent reserve command tool	79
Using ESS or SAN Volume Controller devices directly	81
Using ESS or SAN Volume Controller devices through AIX LVM	82
Migrating a non-SDD volume group to an ESS or SAN Volume Controller SDD multipath volume group in concurrent mode	83
Using the trace function	86
Error log messages	87
Chapter 3. Using SDD on a Hewlett Packard host system	89
Verifying the hardware and software requirements	89
Hardware	89
Software	89
Unsupported environments	89
SDD support for 32-bit and 64-bit applications on HP-UX 11.0 and HP-UX 11i (Fibre Channel)	90
Understanding how SDD works on an HP host system	90
Preparing for SDD installation	90
Configuring the ESS	90
Configuring the SAN Volume Controller	90
Planning for installation	90
Determining if the SDD 1.3.1.5 (or later) server for Expert is installed	91
Installing SDD	92
Postinstallation	94
Upgrading SDD.	96
Upgrading SDD 1.3.0.2 (or earlier) to SDD 1.4.0.0 (or later)	96
Uninstalling the Subsystem Device Driver	97
Configuring SDD	98
Changing an SDD hardware configuration	98
Converting a volume group	98
Dynamic reconfiguration	98
Dynamically changing the SDD path-selection policy algorithm	99
SAN Volume Controller Preferred Node path-selection algorithm	100
SDD datapath query adapter command changes for SDD 1.4.0.0	100
SDD datapath query device command changes for SDD 1.4.0.0	101
SDD server daemon	102
Verifying if the SDD server has started.	102
Starting the SDD server manually	102
Changing to a different port number for the SDD server	102
Stopping the SDD server.	103
Understanding the SDD 1.3.1.5 (or later) support for single-path configuration for ESS	103
Understanding the SDD error recovery policy	103
Single-path mode (for ESS only)	104
Multipath mode	104
Dynamically opening an Invalid or CLOSE_DEAD path.	105
How to import and export volume groups	105
Exporting volume groups.	105
Moving the map file.	106
Creating the volume group device directory	106
Creating the group special file	106
Importing volume groups	106
Using applications with SDD	107
Standard UNIX applications.	107
Installing SDD on a Network File System file server	112

Oracle	113
Chapter 4. Using SDD on a Linux host system.	117
Verifying the hardware and software requirements	117
Hardware	117
Software	117
Unsupported environments	117
Preparing for SDD installation	118
Configuring the ESS	118
Configuring the SAN Volume Controller	118
Configuring Fibre Channel adapters for ESS	118
Configuring Fibre Channel adapters for SAN Volume Controller	118
Installing SDD	119
Upgrading the SDD	119
Verifying the SDD installation	120
Configuring SDD	120
Using multiple commands for SDD configuration	121
Using a single command for SDD configuration	125
Configuring SDD at system startup	125
Maintaining vpath device configuration persistence	125
Dynamically changing the SDD path-selection policy algorithm	126
Dynamically adding paths to SDD devices of a volume group	127
Dynamic reconfiguration	127
Removing SDD	127
SDD server daemon	128
Verifying if the SDD server has started.	128
Starting the SDD server manually	128
Changing to a different port number for the SDD server	128
Stopping the SDD server.	129
Understanding the SDD error recovery policy	129
Understanding the SDD support for single-path configuration	130
Partitioning SDD vpath devices	130
Using standard UNIX applications	131
Chapter 5. Using SDD on a Novell host system	133
Verifying the hardware and software requirements	133
Hardware requirements	133
Software requirements	133
Supported environments	133
Unsupported environments	134
ESS requirements	134
SCSI requirements	134
Fibre Channel requirements.	134
Preparing for SDD installation	135
Configuring the ESS	135
Configuring Fibre Channel adapters.	135
Configuring SCSI adapters	136
Installing SDD.	136
Configuring SDD.	136
Displaying the current version of the SDD	136
Features	136
Automatic path detection, failover and selection	137
Manual operations using the datapath commands	137
Understanding SDD error recovery algorithms	138
Dynamic load balancing	139
ESS logical unit detection	139

Error reporting and logging	139
SDD in NetWare-layered architecture	139
Removing SDD	140
Cluster setup for Novell Network 5.1	140
Cluster setup for Novel Network 6.0	140
Chapter 6. Using SDD on a Sun host system	141
Verifying the hardware and software requirements	141
Hardware	141
Software	141
Supported environments	141
Unsupported environments	142
Understanding how SDD works on a Sun host system	142
Preparing for SDD installation	142
Configuring the ESS	142
Configuring the SAN Volume Controller	142
Determining if the SDD server for Expert is installed.	142
Planning for installation	143
Installing SDD.	144
Verifying the SDD installation	147
Uninstalling SDD.	147
Postinstallation	148
Upgrading the SDD.	150
Configuring SDD.	150
Changing an SDD hardware configuration	150
Dynamically changing the SDD path-selection policy algorithm	152
SAN Volume Controller Preferred Node path-selection algorithm	153
Understanding the SDD 1.3.1.0 (or later) support for single-path configuration for ESS	153
Understanding the SDD error recovery policy	153
SDD server daemon	155
Verifying if the SDD server has started.	155
Starting the SDD server manually	156
Changing to a different port number for the SDD server	156
Stopping the SDD server.	156
Using applications with SDD	157
Standard UNIX applications.	157
Installing SDD on a Network File System file server	157
Oracle	158
Solstice DiskSuite	162
Chapter 7. Using SDD on a Windows NT host system	165
Verifying the hardware and software requirements	165
Hardware	165
Software.	165
Unsupported environments	165
ESS requirements	165
Host system requirements	166
Preparing for SDD installation	166
Configuring the ESS	167
Configuring the SAN Volume Controller	167
Configuring Fibre Channel adapters.	167
Configuring SCSI adapters for ESS devices.	167
Installing SDD.	167
Configuring SDD.	168
Adding paths to SDD devices	168

SAN Volume Controller Preferred Node path-selection algorithm	173
Upgrading the SDD	174
Adding or modifying a multipath storage configuration to the ESS or SAN Volume Controller	174
Removing SDD	176
SDD server daemon	177
Verifying that the SDD server has started.	177
Starting the SDD server manually	177
Changing to a different port number for the SDD server	177
Stopping the SDD server.	178
Displaying the current version of SDD	178
Error recovery and retry policy.	178
Using high availability clustering on an ESS.	180
Special considerations in the high availability clustering environment.	180
Configuring a Windows NT cluster with SDD installed	180
Chapter 8. Using SDD on a Windows 2000 host system	183
Verifying the hardware and software requirements	183
Unsupported environments	183
ESS requirements	183
SAN Volume Controller requirements	183
Host system requirements	184
Preparing for SDD 1.4.0.0 (or later) installation.	184
Configuring the ESS or SAN Volume Controller	185
Configuring Fibre Channel adapters.	185
Configuring SCSI adapters for ESS devices.	185
Installing SDD 1.4.0.0 (or later)	186
Displaying the current version of SDD	188
Upgrading SDD	188
Upgrading to SDD 1.4.0.0 (or later) in a two-node cluster environment	188
Uninstalling SDD.	188
Removing SDD in a two-node cluster environment	189
SDD server daemon	189
Verifying if the SDD server has started.	189
Starting the SDD server manually	190
Changing to a different port number for the SDD server	190
Stopping the SDD server.	190
Adding paths to SDD devices	190
Activating additional paths	191
Verifying that additional paths are installed correctly.	192
SAN Volume Controller Preferred Node path-selection algorithm	193
Error recovery and retry policy.	194
Support for Windows 2000 clustering	195
Special considerations in the Windows 2000 clustering environment	195
Configuring a Windows 2000 cluster with SDD installed	196
Chapter 9. Using SDD on a Windows Server 2003 host system	199
Verifying the hardware and software requirements	199
Unsupported environments	199
ESS requirements	199
Host system requirements	199
Preparing for SDD 1.4.0.0 (or later) installation.	200
Configuring the ESS	201
Configuring Fibre Channel adapters.	201
Configuring SCSI adapters	201
Installing SDD 1.4.0.0 (or later)	201

Displaying the current version of SDD	203
Upgrading SDD	204
Upgrading from a Windows NT host system to Windows Server 2003	204
Uninstalling SDD	204
Removing SDD in a two-node cluster environment	205
SDD server daemon	205
Verifying if the SDD server has started.	205
Starting the SDD server manually	205
Changing to a different port number for the SDD server	206
Stopping the SDD server.	206
Adding paths to SDD devices	206
Activating additional paths	207
Verifying that additional paths are installed correctly.	208
Error recovery and retry policy.	209
Support for Windows Server 2003 clustering	210
Special considerations in the Windows Server 2003 clustering environment	211
Configuring a Windows Server 2003 cluster on ESS.	211
Configure Windows 2003 cluster with SDD installed	211
Chapter 10. The SDD server daemon	215
Understanding how the SDD server daemon works	215
Path reclamation.	215
Path probing	215
Sddsrv and StorWatch Expert	215
Chapter 11. Using the datapath commands	217
datapath query adapter	218
datapath query adaptstats	219
datapath query device.	220
datapath query devstats	222
datapath set adapter	224
datapath set path	225
datapath open path	226
datapath set device policy	228
Notices	229
Trademarks.	230
Electronic emission notices	231
Federal Communications Commission (FCC) statement	231
Industry Canada compliance statement	231
European community compliance statement.	232
Japanese Voluntary Control Council for Interference (VCCI) class A statement	232
Korean government Ministry of Communication (MOC) statement.	233
Chinese Class A warning statement.	233
Taiwan class A compliance statement	233
IBM agreement for licensed internal code.	233
Actions you must not take	234
Glossary	235
Index	255

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.	IBMsdd Driver 32-bit	93
4.	IBMsdd Driver 64-bit	93
5.	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	187
6.	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	203

Tables

1. Package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0	xvi
2. Publications in the SAN Volume Controller library	xix
3. SDD in the protocol stack	1
4. Package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0	7
5. AIX required fixes	8
6. SDD 1.4.0.0 installation packages for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, AIX 5.1.0, and AIX 5.2.0.	15
7. Major files included in the SDD installation package	17
8. Maximum SDD device configuration for ESS LUNs	22
9. PTFs for APARs on AIX with Fibre Channel support and the SDD server daemon running	39
10. Software support for HACMP in concurrent and nonconcurrent mode 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)	43
11. Software support for HACMP in concurrent mode on AIX 5.1.0 (32-bit or 64-bit kernel)	44
12. Software support for HACMP in nonconcurrent mode on AIX 5.1.0 (32-bit and 64-bit kernel)	45
13. Software support for HACMP in concurrent mode	46
14. Software support for HACMP in nonconcurrent mode	47
15. Software support for HACMP/ES on 2105 Model 800	47
16. HACMP and supported SDD features	48
17. APARs on AIX with HACMP and the SDD server daemon running.	48
18. SDD-specific SMIT panels and how to proceed.	69
19. SDD installation scenarios	91
20. Patches necessary for proper operation of SDD on HP-UX 11.0	91
21. SDD components installed for HP host systems	94
22. System files updated for HP host systems	94
23. SDD commands and their descriptions for HP host systems	95
24. SDD components for a Linux host system	120
25. Summary of SDD commands for a Linux host system.	121
26. SDD installation scenarios	143
27. Operating systems and SDD package file names	144
28. SDD components installed for Sun host systems	148
29. System files updated for Sun host systems.	149
30. SDD commands and their descriptions for Sun host systems	149
31. Commands	217

About this book

The IBM[®] Subsystem Device Driver (SDD) provides the multipath configuration environment support for a host system that is attached to an IBM TotalStorage[™] Enterprise Storage Server[®] (ESS) or to an IBM TotalStorage SAN Volume Controller. This book provides step-by-step procedures on how to install, configure, and use SDD for the following host systems:

- IBM AIX[®]
- HP
- Linux Red Hat and SuSE (Intel)
- Novell (ESS only)
- Sun
- 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).

Special characters conventions

The following special character conventions are used in this book:

*** asterisks**

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

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

... ellipsis

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

< > Indicate optional parameters.

Summary of changes

This book contains both information previously presented in the Third Edition (February 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.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

New information

This edition includes the following new information:

- Discussion in this edition supports the ESS and the SAN Volume Controller unless the sections are clearly labeled to support one or the other.
- Support for the multipath configuration environments in ESS devices or SAN Volume Controller devices or both.
- Support for Windows Server 2003 is described in Chapter 9, “Using SDD on a Windows Server 2003 host system”, on page 199.

Modified information

In addition to support of the SAN Volume Controller, the following information has been modified from the previous ESS-only version of this guide:

Chapter 2, “Using SDD on an AIX host system”, on page 7

- Installation package names have been changed as shown in the following table:

Table 1. Package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0

SDD 1.3.3.x	SDD 1.4.0.0	Notes
ibmSdd_432.rte	N/A	Obsolete. This package has been merged with devices.sdd.43.rte
ibmSdd_433.rte	devices.sdd.43.rte	
ibmSdd_510.rte	N/A	Obsolete. This package has been merged with devices.sdd.51.rte
ibmSdd_510nchacmp.rte	devices.sdd.51.rte	
	devices.sdd.52.rte	New package for AIX 5.2.0

- Addition to AIX required fixes
- Corrections

Chapter 4, “Using SDD on a Linux host system”, on page 117 Changes to Linux kernel support include:

- The Redhat 7.x kernel-2.4.18-24.7.xbigmem and kernel-2.4.18-27.7.xbigmem are supported.
- Red Hat Linux Advanced Server 2.1 with:
 - 2.4.9-e.12enterprise
 - 2.4.9-e.16enterprise
- Red Hat Linux Advanced Server 2.1 with the 2.4.9-e.16summit kernel
- Support for SuSE SLES 7 kernel k_smp-2.4.18-224.

- Support for SLES8 kernel k_smp-2.4.19-233 in i686.

Chapter 6, “Using SDD on a Sun host system”, on page 141

- Dynamic reconfiguration.
- Dynamic path-selection algorithms.

Chapter 8, “Using SDD on a Windows 2000 host system”, on page 183

- Removal of the Start Type Configuration Utility requirement.

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
- Other IBM publications that relate to the ESS
- Non-IBM publications that relate to the ESS

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

<http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/2105>

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

Title	Description	Order Number
<i>IBM TotalStorage Enterprise Storage Server: Copy Services Command-Line Interface Reference</i>	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 FlashCopy® configuration tasks within an ESS Copy Services server group.	SC26-7494 (See Note.)
<i>IBM TotalStorage Enterprise Storage Server: Configuration Planner for Open-Systems Hosts</i>	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.)

Title	Description	Order Number
<i>IBM TotalStorage Enterprise Storage Server: Configuration Planner for S/390® and IBM @server zSeries Hosts</i>	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.)
<i>IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide</i>	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.)
<i>IBM TotalStorage Enterprise Storage Server: DFSMS Software Support Reference</i>	This publication provides an overview of the ESS and highlights its unique capabilities. It also describes Data Facility Storage Management Subsystems (DFSMS) software support for the ESS, including support for large volumes.	SC26-7440 (See Note.)
<i>IBM TotalStorage™ Enterprise Storage Server™: Introduction and Planning Guide</i>	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
<i>IBM TotalStorage Enterprise Storage Server: Quick Configuration Guide</i>	This booklet provides flow charts for using the TotalStorage Enterprise Storage Server Specialist (ESS Specialist). The flow charts provide a high-level view of the tasks that the IBM service support representative performs during initial logical configuration. You can also use the flow charts for tasks that you might perform when you are modifying the logical configuration.	SC26-7354
<i>IBM TotalStorage Enterprise Storage Server: System/390® Command Reference</i>	This publication describes the functions of the ESS and provides reference information, such as channel commands, sense bytes, and error recovery procedures for IBM S/390 and zSeries hosts.	SC26-7298
<i>IBM TotalStorage Storage Solutions Safety Notices</i>	This publication provides translations of the danger notices and caution notices that IBM uses in ESS publications.	GC26-7229
<i>IBM TotalStorage Enterprise Storage Server: SCSI Command Reference</i>	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
<i>IBM TotalStorage Enterprise Storage Server: Subsystem Device Driver User's Guide</i>	<p>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.</p> <p>For information about SDD, go to the following Web site: http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/sdd/</p>	SC26-7478
<i>IBM TotalStorage Enterprise Storage Server: User's Guide</i>	This guide provides instructions for setting up and operating the ESS and for analyzing problems.	SC26-7445 (See Note.)
<i>IBM TotalStorage Enterprise Storage Server: Web Interface User's Guide</i>	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
<i>IBM TotalStorage Common Information Model Agent for the Enterprise Storage Server Installation and Configuration Guide</i>	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
<p>Note: No hardcopy book is produced for this publication. However, a PDF file is available from the following Web site:</p> <p>http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/2105/</p>		

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 2. Publications in the SAN Volume Controller library

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

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 xxvi for information about ordering these and other IBM publications.

Title	Description	Order Number
Data-copy services		

Title	Description	Order Number
<i>DFSMS/MVS® Version 1: Advanced Copy Services</i>	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
<i>DFSMS/MVS V1: Remote Copy Guide and Reference</i>	This publication provides guidelines for using remote copy functions with S/390 and zSeries hosts.	SC35-0169
<i>IBM Enterprise Storage Server</i>	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
<i>Implementing Copy Services On S/390</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an 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
<i>IBM TotalStorage ESS Implementing Copy Services in an Open Environment</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure 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		
<i>Fibre Channel Connection (FICON) I/O Interface: Physical Layer</i>	This publication provides information about the fiber-channel I/O interface. This book is also available as a PDF file from the following Web site: www.ibm.com/servers/resourcelink/	SA24-7172
<i>Fibre-channel Subsystem Installation Guide</i>	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
<i>Fibre Transport Services (FTS): Physical and Configuration Planning Guide</i>	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234
<i>IBM SAN Fibre Channel Managed Hub: 3534 Service Guide</i>	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
<i>IBM SAN Fibre Channel Managed Hub: 3534 User's Guide</i>	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

Title	Description	Order Number
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide</i>	<p>This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.</p> <p>To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/</p>	SC26-7350
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide</i>	<p>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.</p> <p>To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/</p>	SC26-7349
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide</i>	<p>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.</p> <p>To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/</p>	SC26-7352
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide</i>	<p>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.</p> <p>To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/</p>	SC26-7351
<i>Implementing Fibre Channel Attachment on the ESS</i>	<p>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.</p>	SG24-6113
Open-systems hosts		
<i>ESA/390: ESCON I/O Interface</i>	<p>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.</p>	SA22-7202
<i>ESS Solutions for Open Systems Storage: Compaq AlphaServer, HP, and Sun</i>	<p>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.</p>	SG24-6119
<i>Fibre-channel Subsystem Installation Guide</i>	<p>This publication tells you how to attach the IBM xSeries 430 and NUMA-Q host systems with fibre-channel adapters to the ESS.</p> <p>Contact your sales representative to obtain this publication.</p>	No order number

Title	Description	Order Number
<i>IBM TotalStorage ESS Implementing Copy Services in an Open Environment</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure 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
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the 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		
<i>Device Support Facilities: User's Guide and Reference</i>	This publication describes the IBM Device Support Facilities (ICKDSF) product that are used with IBM direct access storage device (DASD) subsystems. ICKDSF is a program that you can use to perform functions that are needed for the installation, the use, and the maintenance of IBM DASD. You can also use it to perform service functions, error detection, and media maintenance.	GC35-0033
<i>DFSMS/MVS Version 1: Advanced Copy Services</i>	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
<i>DFSMS/MVS V1: Remote Copy Guide and Reference</i>	This publication provides guidelines for using remote copy functions with S/390 and zSeries hosts.	SC35-0169
<i>ESA/390: ESCON I/O Interface</i>	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
<i>Fibre Transport Services (FTS): Physical and Configuration Planning Guide</i>	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234
<i>Implementing ESS Copy Services on S/390</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an 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

Title	Description	Order Number
<i>ES/9000, ES/3090: IOCP User Guide Volume A04</i>	This publication describes the Input/Output Configuration Program that supports the Enterprise Systems Connection (ESCON) architecture. It describes how to define, install, and configure the channels or channel paths, control units, and I/O devices on the ES/9000 processors and the IBM ES/3090 Processor Complex.	GC38-0097
<i>IOCP User's Guide, IBM @server zSeries</i>	This publication describes the Input/Output Configuration Program that supports the zSeries server. This publication is available in PDF format by accessing ResourceLink at the following Web site: www.ibm.com/servers/resourcelink/	SB10-7029
<i>S/390: Input/Output Configuration Program User's Guide and ESCON Channel-to-Channel Reference</i>	This publication describes the Input/Output Configuration Program that supports ESCON architecture and the ESCON multiple image facility.	GC38-0401
<i>IBM z/OS Hardware Configuration Definition User's Guide</i>	This guide provides conceptual and procedural information to help you use the z/OS Hardware Configuration Definition (HCD) application. It also explains: <ul style="list-style-type: none"> • How to migrate existing IOCP/MVSCP definitions • How to use HCD to dynamically activate a new configuration • How to resolve problems in conjunction with MVS/ESA HCD 	SC33-7988
<i>OS/390: Hardware Configuration Definition User's Guide</i>	This guide provides detailed information about the input/output definition file and about how to configure parallel access volumes. This guide discusses how to use Hardware Configuration Definition for both OS/390® and z/OS V1R1.	SC28-1848
<i>OS/390 V2R10.0: MVS System Messages Volume 1 (ABA - ASA)</i>	This publication lists OS/390 MVS™ system messages ABA to ASA.	GC28-1784
<i>Using IBM 3390 Direct Access Storage in a VM Environment</i>	This publication provides device-specific information for the various models of the 3390 and describes methods you can use to manage storage efficiently using the VM operating system. It provides guidance on managing system performance, availability, and space through effective use of the direct access storage subsystem.	GG26-4575
<i>Using IBM 3390 Direct Access Storage in a VSE Environment</i>	This publication helps you use the 3390 in a VSE environment. It includes planning information for adding new 3390 units and instructions for installing devices, migrating data, and performing ongoing storage management activities.	GC26-4576
<i>Using IBM 3390 Direct Access Storage in an MVS Environment</i>	This publication helps you use the 3390 in an MVS environment. It includes device-specific information for the various models of the 3390 and illustrates techniques for more efficient storage management. It also offers guidance on managing system performance, availability, and space utilization through effective use of the direct access storage subsystem.	GC26-4574
<i>z/Architecture Principles of Operation</i>	This publication provides a detailed definition of the z/Architecture™. It is written as a reference for use primarily by assembler language programmers and describes each function at the level of detail needed to prepare an assembler language program that relies on a particular function. However, anyone concerned with the functional details of z/Architecture will find this publication useful.	SA22-7832
SAN		

Title	Description	Order Number
<i>IBM OS/390 Hardware Configuration Definition User's Guide</i>	<p>This guide explains how to use the Hardware Configuration Data application to perform the following tasks:</p> <ul style="list-style-type: none"> • Define new hardware configurations • View and modify existing hardware configurations • Activate configurations • Query supported hardware • Maintain input/output definition files (IODFs) • Compare two IODFs or compare an IODF with an actual configuration • Print reports of configurations • Create graphical reports of a configuration • Migrate existing configuration data 	SC28-1848
<i>IBM SAN Fibre Channel Managed Hub: 3534 Service Guide</i>	<p>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.</p>	SY27-7616
<i>IBM SAN Fibre Channel Managed Hub: 3534 User's Guide</i>	<p>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.</p>	GC26-7391
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide</i>	<p>This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.</p> <p>To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/</p>	SC26-7350
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide</i>	<p>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).</p> <p>To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/</p>	SC26-7349
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide</i>	<p>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.</p> <p>To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/</p>	SC26-7352
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide</i>	<p>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.</p> <p>To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/</p>	SC26-7351

Title	Description	Order Number
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the 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
Seascope family		
<i>IBM Enterprise Storage Server</i>	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
<i>IBM Enterprise Storage Server Performance Monitoring and Tuning Guide</i>	This guide, from the IBM International Technical Support Organization, provides guidance on the best way to configure, monitor, and manage your ESS to ensure optimum performance.	SG24-5656
<i>IBM Versatile Storage Server: Introduction and Planning Guide</i>	This publication introduces the IBM Versatile Storage Server™ and lists the features you can order. It also provides planning information for both 2105 Models B09 and 100.	GC26-7223
<i>Implementing the IBM Enterprise Storage Server in Your Environment</i>	This publication, from the IBM International Technical Support Organization, can help you install, tailor, and configure the ESS in your environment.	SG24-5420
Storage management		
<i>Device Support Facilities: User's Guide and Reference</i>	This publication describes the IBM Device Support Facilities (ICKDSF) product used with IBM direct access storage device (DASD) subsystems. ICKDSF is a program that you can use to perform functions that are needed for the installation, the use, and the maintenance of IBM DASD. You can also use it to perform service functions, error detection, and media maintenance.	GC35-0033
<i>IBM TotalStorage Solutions Handbook</i>	This handbook, from the IBM International Technical Support Organization, helps you understand what makes up enterprise storage management. The concepts include the key technologies that you must know and the IBM subsystems, software, and solutions that are available today. It also provides guidelines for implementing various enterprise storage administration tasks so that you can establish your own enterprise storage management environment.	SG24-5250
<i>IBM TotalStorage Expert: Hands-On Usage Guide</i>	This guide, from the IBM International Technical Support Organization, helps you install, tailor, configure, and use TotalStorage ESS Expert.	SG24-6102
<i>IBM TotalStorage Expert Installation Guide</i>	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
<i>Using IBM 3390 Direct Access Storage in a VM Environment</i>	This publication provides device-specific information for the various models of the 3390 and describes methods 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

Title	Description	Order Number
<i>Using IBM 3390 Direct Access Storage in a VSE Environment</i>	This publication helps you use the 3390 in a VSE environment. It includes planning information for adding new 3390 units and instructions for installing devices, migrating data, and performing ongoing storage management activities.	GC26-4576
<i>Using IBM 3390 Direct Access Storage in an MVS Environment</i>	This publication helps you use the 3390 in an MVS environment. It includes device-specific information for the various models of the 3390 and illustrates techniques for more efficient storage management. It also offers guidance 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
<i>Quick Start Guide: An Example with Network File System (NFS)</i>	This guide tells you how to configure the Veritas Cluster Server. See also the companion document, <i>Veritas Cluster Server User's Guide</i> .
<i>Veritas Cluster Server Installation Guide</i>	This guide tells you how to install the Veritas Cluster Server. See also the companion document, <i>Veritas Cluster Server Release Notes</i> .
<i>Veritas Cluster Server Release Notes</i>	These release notes tell you how to install the Veritas Cluster Server. See also the companion document, <i>Veritas Cluster Server Installation Guide</i> .

Title	Description
<i>Veritas Cluster Server User's Guide</i>	This guide tells you how to configure the Veritas Cluster Server. See also the companion document, <i>Quick Start Guide: An Example with Network File System (NFS)</i> .
<i>Veritas Volume Manager Hardware Notes</i>	These hardware notes tell you how to implement multiple paths dynamically.
<i>Veritas Volume Manager Installation Guide</i>	This guide tells you how to install VxVM.
<i>Veritas Volume Manager Storage Administrators Guide</i>	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 or any other ESS documentation, you can submit them in one of the following ways:

- E-mail

Submit your comments electronically to the following e-mail address:
 starpubs@us.ibm.com

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

- Mail 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
RCF Processing Department
M86/050
5600 Cottle Road
San Jose, CA 95193-0001
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

SDD resides above the disk driver of a host system in the protocol stack and acts as a pseudo device driver. See the diagrams in Table 3. 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 dynamically switches to another path.

Table 3. SDD in the protocol stack




Table 3. SDD in the protocol stack (continued)

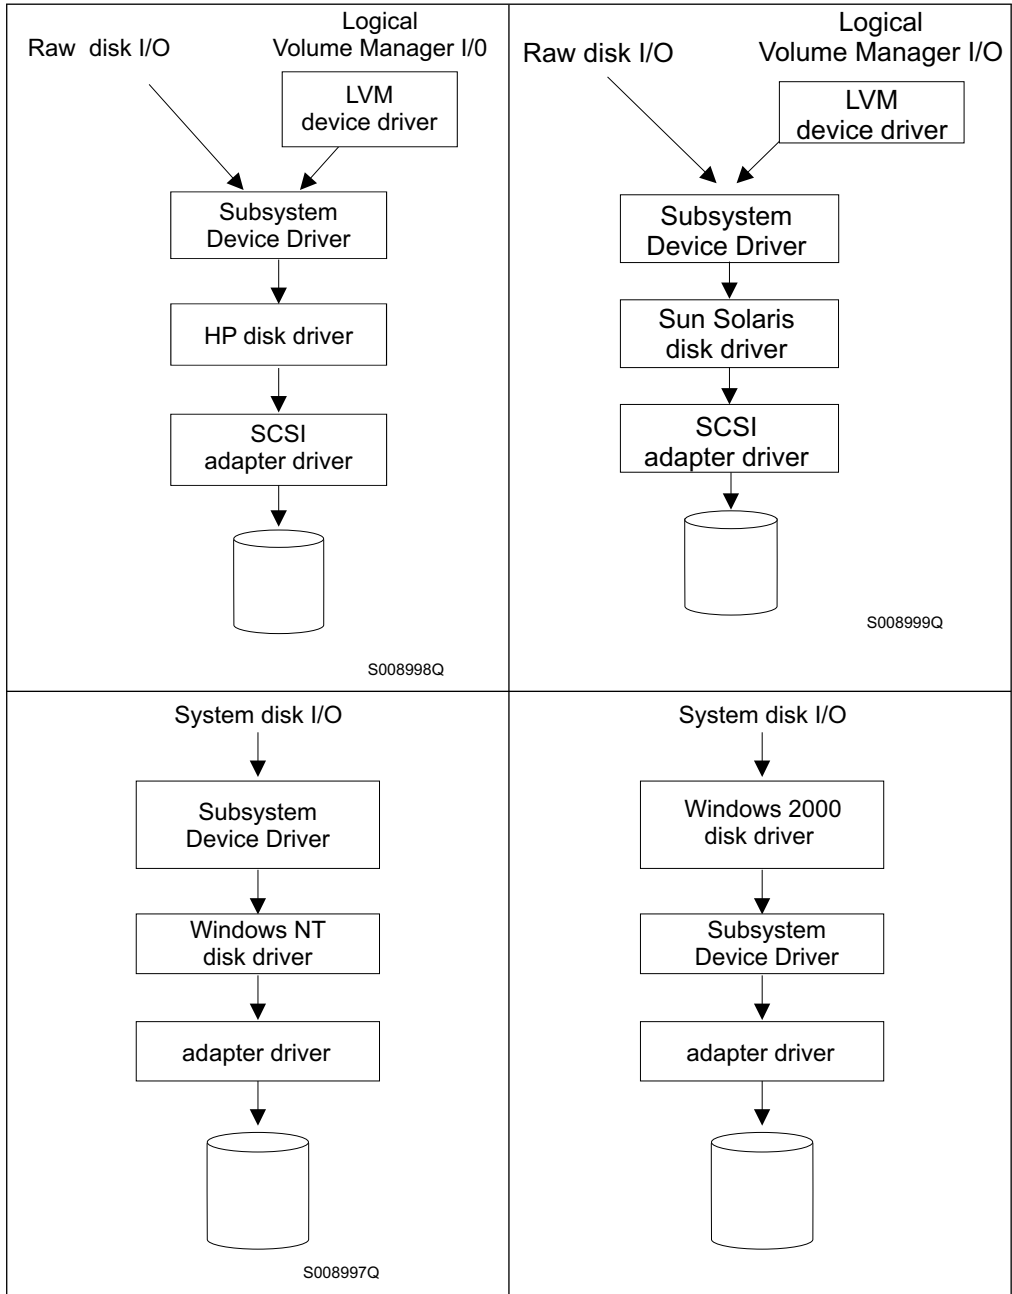
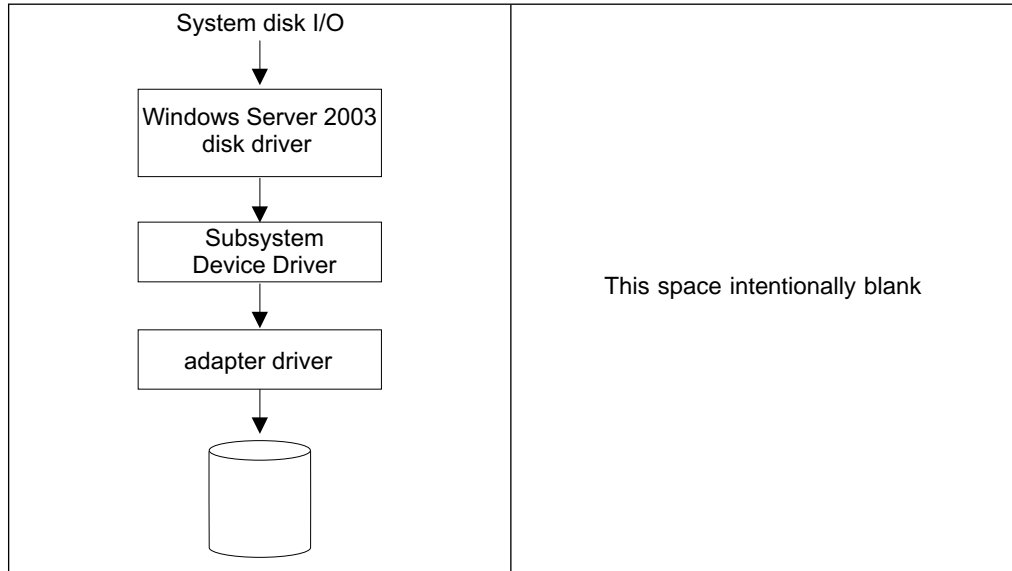


Table 3. SDD in the protocol stack (continued)



Each SDD device represents a unique physical device on the storage server. 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 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.

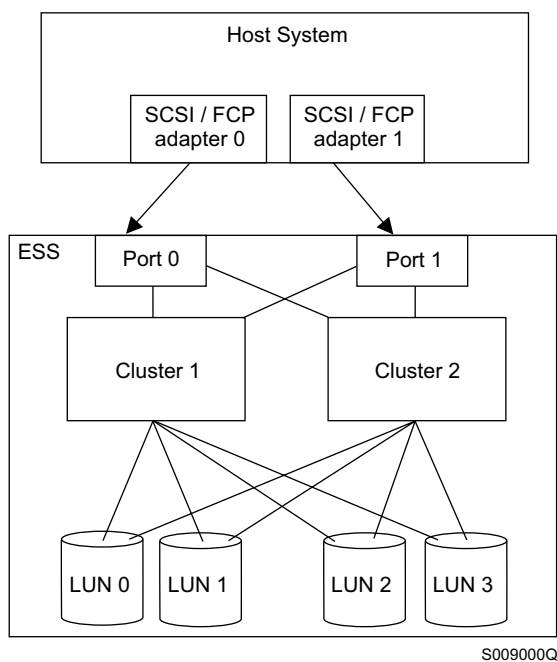


Figure 1. Multipath connections between a host system and the disk storage in an ESS

Figure 2 on page 5 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.

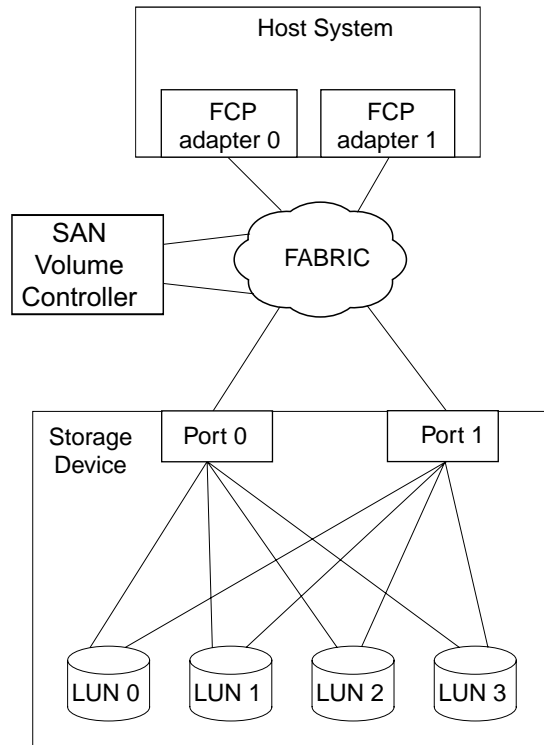


Figure 2. Multipath connections between a host system and the disk storage with a SAN Volume Controller

Note: The SAN Volume Controller does not support 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. During the download and installation process, the host adapters inside the ESS might not respond to host I/O requests for approximately 30 seconds.

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.

To perform concurrent download of LIC for the 2105 Models E10/E20 F10/F20, refer to the *2105 Model E10/E20 F10/F20 Microcode Installation Instructions* for instructions.

To perform concurrent download of LIC for the 2105 Model 800, refer to the *2105 Model 800 Microcode Installation Instructions* for instructions.

Chapter 2. Using SDD on an AIX host system

This chapter provides step-by-step procedures on how to install, configure, upgrade, and remove SDD on an AIX host system that is attached to an ESS or to a SAN Volume Controller. SDD 1.4.0.0 supports the coexistence of ESS devices and SAN Volume Controller devices. For updated and additional information that is not included in this chapter, see the README file on the compact disc or visit the SDD Web site:

www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

Verifying the hardware and software requirements

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

Hardware

- ESS or SAN Volume Controller or both
- A switch if using a SAN Volume Controller (no direct attachment allowed for SAN Volume Controller)
- Host system
- SCSI adapters and cables
- Fibre 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

Package names have changed from all previous package names in order to comply with AIX packaging rules and allow for NIM installation. Table 4 shows the package-naming relationship between SDD 1.3.3.x and 1.4.0.0.

Table 4. Package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0

SDD 1.3.3.x	SDD 1.4.0.0	Notes
<code>ibmSdd_432.rte</code>	N/A	Obsolete. This package has been merged with <code>devices.sdd.43.rte</code> .
<code>ibmSdd_433.rte</code>	<code>devices.sdd.43.rte</code>	
<code>ibmSdd_510.rte</code>	N/A	Obsolete. This package has been merged with <code>devices.sdd.51.rte</code> .
<code>ibmSdd_510nchacmp.rte</code>	<code>devices.sdd.51.rte</code>	
	<code>devices.sdd.52.rte</code>	New package for AIX 5.2.0 (and up).

Note: SDD 1.4.0.0 (and 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.

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 4.3.2, 4.3.3, and 5.1.0
- 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 you are using AIX 4.3.2, AIX 4.3.3 or AIX 5.1.0.

Host system requirements

To successfully install SDD for AIX, you must have AIX 4.3.2, AIX 4.3.3, AIX 5.1.0, or AIX 5.2.0 installed on your host system along with the program temporary fixes (PTFs) in Table 5.

To successfully install SDD for SAN Volume Controller, you must have AIX 4.3.3, AIX 5.1.0, or AIX 5.2.0 installed on your host system along with the PTFs shown in Table 5.

Table 5. AIX required fixes

AIX level	PTF number	Component name	Component level
4.3.2	U461953	bos.rte.lvm	4.3.2.4
4.3.3	U483803	sysxfcp	4.3.3.84
4.3.3.	U484723	sysxfcp	4.3.3.86
5.1.0	U476971	sysxfcp	5.1.0.28
5.1.0	U482718	sysxfcp	5.1.0.35
5.1.0	U483680	sysxfcp	5.1.0.36
5.1.F	U483641	sysxlvms (bos.rte.lvm)	5.1.0.37
5.1.G	U484105	sysxlvms (bos.rte.lvm)	5.1.0.50
5.2.0	U486795	sysxlvms (bos.rte.lvm)	5.2.0.2
5.2.B	U485120	sysxlvms (bos.rte.lvm)	5.2.0.10

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

techsupport.services.ibm.com/server/support

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:

- IBM TotalStorage SAN Volume Controller 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 System Modification Program (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:

techsupport.services.ibm.com/server/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:

techsupport.services.ibm.com/server/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*.

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

www.ibm.com/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® / 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.

For information about the Fibre Channel adapters that can be used on your AIX host system go to the following Web site:

storage.ibm.com/hardsoft/products/ess/supserver.htm

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
- 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

Notes:

1. SDD 1.3.3.9 (or later) supports manual exclusion of ESS devices from the SDD configuration. If you want to manually exclude ESS devices from the SDD configuration, you must do so before configuring SDD devices.
2. SDD 1.4.0.0 (or later) supports manual exclusion of SAN Volume Controller devices from the SDD configuration. If you want to manually exclude SAN Volume Controller devices from the SDD configuration, you must do so before configuring SDD devices.
3. The **querysn** command can be used to exclude any ESS devices from the SDD configuration. The **querysn** command reads the unique serial number of an ESS device (hdisk) and saves the serial number in an exclude file. See “Manual exclusion of ESS devices from the SDD configuration” on page 16.
4. With SDD 1.4.0.0 (or later), the **querysn** command can also be used to exclude any SAN Volume Controller devices from the SDD configuration. The **querysn** command reads the unique serial number of a SAN Volume Controller device (hdisk) and saves the serial number in an exclude file.

See “querysn” on page 78 for additional information about the **querysn** command.

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.

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:

techsupport.services.ibm.com/server/fixes

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, type `smitty install_update` and press Enter to go directly to the installation panels. The **Install and Update Software** menu displays.
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 displays.
8. Highlight **Software to Install** and press **F4**. The **Software to Install** panel displays.
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.

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

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

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

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

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

where **nnnnn** represents the APAR numbers.

If the APARs are listed, that means that they are installed. If they are installed, go to “Configuring Fibre Channel-attached devices” on page 13. 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. Type `smitty deinstall` at the AIX command prompt and press Enter. The Remove Installed Software panel displays.
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. Type `installp -ug devices.pci.df1000f9` and press Enter.
2. Type `installp -ug devices.pci.df1000f7` and press Enter.
3. Type `installp -ug devices.common.IBM.fc` and press Enter.
4. Type `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” 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 issue the following command for each installed FCP adapter:

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

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

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

Verifying the adapter firmware level

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

techsupport.services.ibm.com/server/mdownload2/download.html#adapter

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. Issue the **lscfg -vl fcsN** command. The vital product data for the adapter displays.
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 13). 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, type `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 displays 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.
8. After the download is complete, issue the `lscfg -v -l fcsN` command to verify the firmware level on each Fibre Channel adapter.

Determining if the `sddServer` for Expert is installed

If you previously installed the 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 `sddServer` 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, type:

```
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.ssddom02.storage.ibm.com/techsup/webnav.nsf/support/swexpert

For more information about the `sddServer` daemon, go to “SDD server daemon” on page 37

Planning for SDD installation on a pSeries 690 server LPAR

As a standard feature, the IBM *e*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 6 to determine the correct installation package.

Determining the installation package

Before you install SDD on your AIX 4.3.2, AIX 4.3.3, AIX 5.1.0, or AIX 5.2.0 host system, you need to determine the installation package that is appropriate for your environment.

Installation packages for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, AIX 5.1.0, and AIX 5.2.0 host systems

Table 6. SDD 1.4.0.0 installation packages for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, AIX 5.1.0, and AIX 5.2.0

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

Note: ¹ SAN Volume Controller devices are not supported on AIX 4.3.2.

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 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 ESS 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 ESS 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 ESS device with multipath configuration. Refer to chapter 6 of the *IBM TotalStorage Host System Attachment Guide* for additional information.

The SDD driver will automatically exclude any ESS devices from SDD configuration if these ESS 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 ESS boot devices, you should use the **querysn** command to manually exclude all ESS 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 ESS devices from the SDD configuration. The **querysn** command reads the unique serial number of an ESS 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 ESS devices from the SDD configuration. See “querysn” on page 78 for the syntax of the **querysn** command.

The exclude file, /etc/vpexclude, holds the serial numbers of all ESS 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.

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. However, if you want to exclude any SAN Volume Controller devices from the SDD configuration, you may still do so with the **querysn** command.

Replacing manually excluded ESS devices in the SDD configuration

Use the following procedure to place manually-excluded ESS 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 ESS 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 ESS devices.
4. Execute **lsvpcfg** to verify that these ESS 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 7 lists the major files that are part of the SDD installation package.

Table 7. 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.
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 ESS hdisk device volume group to a Subsystem Device Driver vpath device volume group.
vp2hd	The SDD script that converts an SDD vpath device volume group to an ESS 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 ESS devices.
mkvg4vp	The command that creates an SDD volume group.
extendvg4vp	The command that extends SDD devices to an SDD volume group.

Table 7. Major files included in the SDD installation package (continued)

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.
sddsrv.conf	The SDD server configuration file.
lvmrecover	The SDD script that restores a system's SDD devices and LVM configuration when a migration failure occurs

Determining the installation type

Before you install SDD on your AIX 4.3.2, AIX 4.3.3, AIX 5.1.0, or AIX 5.2.0 host system, 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” 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 26 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 29 for instructions. A PTF file has the letters bff as its file extension (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 6 on page 15 lists and describes SDD support for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, AIX 5.1.0 and AIX 5.2.0.

Note:

- You must have root access and AIX system administrator knowledge to install SDD.
- The following procedures assume that SDD will be used to access all of your single-path and multipath devices.
- If you are installing an *older* version of SDD, you must first remove any previously installed, *newer* version of SDD from your host system.

Use the System Management Interface Tool (SMIT) facility to install SDD. The SMIT facility has two interfaces, nongraphical (type `smitty` to invoke the nongraphical user interface) and graphical (type `smit` to invoke the graphical user interface).

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, type `smitty install_update` and press Enter to go directly to the installation panels. The Install and Update Software menu displays.
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 displays.
8. Highlight **Software to Install** and press **F4**. The Software to Install panel displays.
9. Select the installation package that is appropriate for your environment.
10. Press Enter. The Install and Update from LATEST Available Software panel displays 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.  
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.

Tip:

For AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0, the published limitation on one system is 10 000 devices. 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 the number of paths to each disk.

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 issuing the following command:

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

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

Alternatively, you can issue 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

Note: 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 1.3.2.9 (or later) is installed. Go to “Verifying if the SDD server has started” on page 37 to see if the SDD server has automatically started. Go to “SDD server daemon” on page 37 for more details about the SDD server daemon.

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 issuing the following command:

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

Alternatively, you can issue 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 **lspp -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 a configuration

For AIX 5.2.0 (and later), you can have a configuration with a maximum of 1000 ESS LUNs. In earlier AIX versions, SDD limits the maximum number of ESS LUNs to 600. In SDD 1.4.0.0, the total LUNs supported by SDD on AIX 4.3.3 and AIX 5.1 is 600. This means that if an AIX host has both ESS and SAN Volume Controller storage attached, the two different types of devices will share the maximum 600 as the total number of devices that can be configured.

For AIX 5.2.0 (and later), you can have a configuration with a maximum of 512 SAN Volume Controller LUNs. A new SDD attribute, **2145_max_luns**, defines the maximum number of SAN Volume Controller LUNs that can be configured on a host. The maximum value for **2145_max_luns** is 512. You cannot change the value of the **2145_max_luns** attribute.

You can have a maximum of 32 paths per vpath. Since the number of paths influences performance, you should use the minimum number of paths necessary to achieve sufficient redundancy in the switching fabric.

Attention: This new maximum number for LUN configuration only applies to the `devices.sdd.52.rte` package.

Maximum SDD device configuration for ESS LUNs

Starting from SDD 1.4.0.0, the maximum SDD device configuration for ESS LUNs and the maximum paths per SDD device is given in Table 8.

Table 8. Maximum SDD device configuration for ESS LUNs

Number of LUNs	Maximum paths per vpath
1- 600 vpath LUN	16
601 - 900 vpath LUN	8
901 - 1000 vpath LUN	4

A new SDD attribute, **2105_max_luns**, defines the maximum number of ESS LUNs that can be configured on a host. The range of valid values for **2105_max_luns** is 600 - 1000 in increments of 100. The default value for **2105_max_luns** is 600.

The system administrator must ensure that the number of paths (hdisks) configured for each LUN does not exceed the maximum number of paths shown in Table 8. If

the number of paths (hdisks) for each LUN does exceed the maximum number of paths shown in Table 8 on page 22, the SDD configuration process will exit without configuring SDD devices.

The **addpaths** command follows the same rules identified in Table 8 on page 22.

This design is based on the current LVM design of PVID on AIX. With extra hdisks that are not configured by SDD, LVM can pick up the wrong physical volume instead of the SDD device during volume group vary on. This will cause the loss of single-point-failure protection provided by SDD.

After installing SDD, use following procedure to configure more than 600 vpaths:

Note: This procedure is for a clean installation of SDD only, where the host did not have a previous version of SDD installed. If you have migrated or upgraded SDD, see “Adding and configuring more ESS LUNs with AIX 5.2.0” on page 27.

1. Execute '**cfgmgr -l [scsiN/fcsN]**' for each relevant SCSI or FCP adapter.

Due to large number of hdisks configured on the AIX host, you should ensure that the AIX file system '/'s size is large enough to hold the device special files. The minimum required size is 8 MB (1 MB equals approximately 1 000 000 bytes). For example:

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

The AIX Fibre Channel adapter has an attribute named `lg_term_dma` that controls the DMA memory resource that an adapter driver can use. When you have more than 600 LUNs configured on an AIX 5.2 system, the device open process may fail due to the lack of DMA memory resource. You should increase the DMA memory from the default value of 0x200000 to 0x400000 if you plan to configure more than 600 LUNs. Changing this Fibre Channel adapter attribute requires that you unconfigure and then reconfigure the Fibre Channel adapter and all its children devices. It is better to change this `lg_term_dma` attribute before you assign or configure ESS LUNs on a host system. Because unconfiguring and reconfiguring a large number of devices is very time-consuming, you should perform the following steps:

- a. Execute '**lsattr -El fcsN**' to check the current value of `lg_term_dma`.
 - b. Redefine all existing Fibre Channel adapters by issuing '**rmdev -l fcsN -R**'. It may take long time to unconfigure a large number of devices. An alternative method to speed up this step is to disconnect all Fibre Channel cables and reboot the system.
 - c. Use '**chdev -l fcsN -a lg_term_dma=0x400000**' to increase the DMA value.
 - d. If you disconnected the Fibre Channel cables in step 1b, then reconnect the cables.
 - e. Assign the new LUNs to the AIX host
 - f. Configure the Fibre Channel adapters, its children devices and hdisks using **cfgmgr -l fcsN**.
2. Execute '`/usr/lib/methods/defdpo`'.
 3. 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 8 on page 22.

4. Determine that no more than the allowed number of hdisks are configured for each LUN. Otherwise, you need to remove the extra hdisks. (if the maximum LUN is intended between 901-1000, then no more than 4 hdisks should be configured).
5. Execute '**cfallvpath**' to configure SDD vpath devices.

Preparing to configure SDD

Before you configure SDD, ensure that:

- The ESS or SAN Volume Controller is operational.
- The `devices.sdd.nn.rte` software is installed on the AIX host system, where *nn* identifies the installation package.
- The ESS or SAN Volume Controller hdisks are configured correctly on the AIX host system.

Configure the ESS devices or SAN Volume Controller devices before you configure SDD. If you configure multiple paths to an ESS device or to a SAN Volume Controller device, make sure that all paths (hdisks) are in Available state. Otherwise, some SDD devices will lose multipath capability.

Perform the following steps:

1. Issue the **lsdev -Cc disk | grep 2105** command to check the ESS device configuration or issue the **lsdev -Cc disk | grep SAN Volume Controller** command to check the SAN Volume Controller device configuration.
2. If you have already created some ESS or SAN Volume Controller volume groups, vary off (deactivate) all active volume groups with ESS or SAN Volume Controller subsystem disks by using the **varyoffvg** (LVM) command.

Attention: Before you vary off a volume group, unmount all file systems in that volume group. If some ESS or SAN Volume Controller devices (hdisks) are used as physical volumes of an active volume group and if there are file systems of that volume group being mounted, you must unmount all file systems and vary off all active volume groups with ESS or SAN Volume Controller SDD disks, in order to configure SDD vpath devices correctly.

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. Type `smitty device` from your desktop window. The Devices menu displays.
2. Highlight **Data Path Device** and press Enter. The Data Path Device panel displays.
3. Highlight **Define and Configure All Data Path Devices** and press Enter. The configuration process begins.
4. Check the SDD configuration state. See "Displaying the ESS or SAN Volume Controller vpath device configuration" on page 60.
5. Enter the **varyonvg** command to vary on all deactivated ESS or SAN Volume Controller volume groups.
6. If you want to convert the ESS or SAN Volume Controller hdisk volume group to SDD vpath devices, you must run the `hd2vp` utility. (See "hd2vp and vp2hd" on page 76 for information about this utility.)
7. Mount the file systems for all volume groups that were previously unmounted.

Unconfiguring Subsystem Device Driver

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.

Then, run the `vp2hd` conversion script to convert the volume group from SDD devices (`vpathN`) to ESS or SAN Volume Controller 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 49.

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 *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 -I vpathN` command to return the device to the Available state.

If you 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 24.

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. Type `smitty device` from your desktop window. The Devices menu displays.
2. Highlight **Devices** and press Enter.
3. Highlight **Data Path Device** and press Enter. The Data Path Device panel displays.
4. Highlight **Remove a Data Path Device** and press Enter. A list of all SDD devices and their states (either Defined or Available) displays.
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 from the Available state to the Defined state is: `rmdev -I dpo -R`. The fast-path command to 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 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 StorWatch 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 change. These activities are not allowed during SDD migration or upgrade (for example, configuring more devices).

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 using `smitty install`. The migration or upgrade script to save the LVM configuration will be executed as part of the install procedure initiated via `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.  
Migration from SDD to SDD is successfully completed!
```

Error recovery for migration or upgrade

If any error occurred during the pre-installation or post-installation, such as disconnection of cables, you can recover the migration or upgrade with the following instructions.

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 install the new SDD package with Smitty install.

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.

Adding and configuring more ESS LUNs with AIX 5.2.0

You can add and configure more ESS LUNs to an AIX 5.2 system and exceed the default value of the **2105_max_luns** attribute. In this case, you will need to change the value of **2105_max_luns** prior to configuring hdisks on the AIX host.

The following minimum host system resource requirements are necessary to configure more than 600 LUNs:

1. Due to large number of hdisks configured on the AIX host, you should ensure that the AIX file system `/`'s size is large enough to hold the device special files. The smallest value you should use is 8 MB (1 MB equals approximately 1 000 000 bytes). For example:

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

2. The AIX Fibre Channel adapter attribute `lg_term_dma` controls the DMA memory resource that an adapter driver can use. When you have more than 600 LUNs configured on an AIX 5.2 system, the device open process may fail because of the lack of the DMA memory resource. You should increase the DMA memory from the default value of `0x200000` to `0x400000` if you plan to configure more than 600 LUNs. Changing this Fibre Channel adapter attribute requires that you unconfigure and then reconfigure the Fibre Channel adapter and all its children devices. It is better to change this `lg_term_dma` attribute before assigning and configuring ESS LUNs on a host system. Before assigning LUNs from the ESS to the AIX host and configuring these assigned LUNs, you should:
 - a. Execute `'lsattr -El fcsN'` to check the current value of `lg_term_dma`.
 - b. Redefine all existing Fibre Channel adapters by issuing `'rmdev -l fcsN -R'`. It may take long time to unconfigure a large number of devices. An alternative method to speed up this step is to disconnect all Fibre Channel cables and reboot the system.
 - c. Use `'chdev -l fcsN -a lg_term_dma=0x400000'` to increase the DMA value.
 - d. If you disconnected the Fibre Channel cables in step 2b, reconnect the cables.

- e. Assign the new LUNs to the AIX host
- f. Configure the Fibre Channel adapters, its children devices and hdisks using **cfgmgr -l fcsN**.

To increase the number of vpaths that SDD will configure on your host system, use the following procedure to change the value of the **2105_max_luns** attribute :

1. 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 command:

```
lsattr -El dpo
persistent_resv yes N/A False
2105_max_luns 600 Data Path Optimizer Pseudo Device Driver True
2145_max_luns 512 Data Path Optimizer Parent False".
```

2. Determine how many paths each vpath currently has by issuing the command **datapath query device**.
3. Change all the existing vpaths state to DEFINED by executing **rmdev -l dpo -R**.
4. 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 8 on page 22 for maximum paths allowed.
5. 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 8 on page 22.
6. Execute **cfallvpath** to configure SDD vpath devices.

Note: If step 3 is omitted, step 5 will fail and the following error message will be displayed:

```
Method error (/usr/lib/methods/chggen):
0514-029 Cannot perform the requested function because a child device of the
specified device is not in a correct state.
```

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 will need to install the corresponding SDD package for that OS level. Use the following procedure:

1. Uninstall the SDD package. See "Removing SDD from an AIX host system" on page 33 for detailed instructions.
2. Upgrade the AIX OS and reboot if rebooting does not occur automatically after the OS upgrade.
3. Install the new SDD package corresponding to the upgraded OS level. See "Installing SDD" on page 18 for detailed instructions.

Important:

1. Migration of an SDD package from a lower OS level to a higher OS level after an OS upgrade is not supported. For example, migration from `devices.sdd.43.rte` to `devices.sdd.51.rte` after and OS upgrade from AIX 4.3.3. to AIX 5.1 is not supported.
2. 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.

3. The maximum supported LUNs for the SAN Volume Controller on AIX 4.3.3 and AIX 5.1 is 600, whereas the maximum supported LUNs for the SAN Volume Controller on AIX 5.2 is 512.

Updating SDD packages by applying a program temporary fix

SDD 1.4.0.0 and later allows you 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 may commit or reject the PTF at a later time, without uninstalling SDD from the host system.

Use the System Management Interface Tool (SMIT) facility to update SDD. The SMIT facility has two interfaces, nongraphical (type `smitty` to invoke the nongraphical user interface) and graphical (type `smit` to invoke the graphical user interface).

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, type `smitty install_update` and press Enter to go directly to the installation panels. The Install and Update Software menu displays.
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 displays.
8. Highlight **Software to Install** and press F4. The Software to Install panel displays.
9. Select the PTF package that you want to install.
10. Press Enter. The Install and Update from LATEST Available Software panel displays with the name of the software that you selected to install.
11. If you only want to apply the PTF, highlight **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

Use the SMIT facility to update SDD. The SMIT facility has two interfaces: nongraphical (type `smitty` to invoke the nongraphical user interface) and graphical (type `smit` to invoke the graphical user interface).

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, type `smitty install` and press Enter to go directly to the installation panels. The Software Installation and Maintenance menu displays.
3. Highlight **Software Maintenance and Utilities** and press Enter.
4. Highlight **Commit Applied Software Updates** to commit the PTF or highlight **Reject Applied Software Updates** to reject the PTF.
5. Press Enter. The Commit Applied Software Updates panel displays or the Reject Applied Software Updates panel displays.
6. Highlight **Software name** and press F4. The software name panel displays.
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 may indicate that a reboot is necessary.

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. Type `smitty device` from your desktop window. The Devices menu displays.
2. Select **Data Path Device** and press Enter. The Data Path Device panel displays.
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 displays.

If any device is listed as Defined, the configuration was not successful. Check the configuration procedure again. See “Configuring SDD” on page 24 for the procedure.

Perform the following steps to verify that multiple paths are configured for *each* adapter connected to an ESS or SAN Volume Controller port:

1. Type `smitty device` from your desktop window. The Devices menu displays.
2. Highlight **Data Path Device** and press Enter. The Data Path Device panel displays.
3. Highlight **Display Data Path Device Adapter Status** and press Enter. All attached paths for each adapter are displayed.

If you want to use the command-line interface to verify the configuration, type `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 LUN (for example, `60056768018A0210B00000000000006B`)
- The names of the AIX disk devices making up the pseudo device and their configuration and physical volume state

Changes to the Object Data Manager (ODM) database

The information in this section concerns developers who intend to use any ODM data provided by SDD. Changes have been made to the ODM database to include:

- Displaying the device model
- Displaying the serial number

Displaying the device model

Prior to SDD 1.4.0.0, the `connwhere` field in `CuDv` object displays the device serial number. In SDD 1.4.0.0, this field displays the device model.

For example, in previous SDD versions:

```
[root@anonymous]/> odmget -q "name=vpath0" CuDv
```

outputs the following information:

```
CuDv:
  name = "vpath0"
  status = 1
  chgstatus = 0
  ddins = "vpathdd"
  location = ""
  parent = "dpo"
  connwhere = "31412028"
  PdDvLn = "disk/dpo/vpath"
```

In SDD 1.4.0.0:

This example shows a customized object for an ESS device:

```
[root@anonymous]> odmget -q "name=vpath0" CuDv
CuDv:
name = "vpath0"
status = 1
chgstatus = 0
ddins = "vpathdd"
location = ""
parent = "dpo"
connwhere = "2105800"
PdDvLn = "disk/dpo/vpath"
```

This example shows a customized object for a SAN Volume Controller device:

```
[root@anonymous]> odmget -q "name=vpath40" CuDv
CuDv:
name = "vpath40"
status = 1
chgstatus = 0
ddins = "vpathdd"
location = ""
parent = "dpo"
connwhere = "2145"
PdDvLn = "disk/dpo/vpath"
```

Displaying the serial number

The following examples display the serial number:

This example shows the `serial_number` attribute for an ESS device:

```
[root@anonymous]> odmget -q "name=vpath0 and attribute=serial_number" CuAt
CuAt:
name = "vpath0"
attribute = "serial_number"
value = "01523922"
type = "R"
generic = "D"
rep = "s"
nls_index = 5
```

This example shows the `serial_number` attribute for a SAN Volume Controller device:

```
[root@anonymous]>odmget -q "name=vpath40 and attribute=serial_number" CuAt
CuAt:
name = "vpath40"
attribute = "serial_number"
value = "6005676801898210B000000000000000"
type = "R"
generic = "D"
rep = "s"
nls_index = 5
```

Removing SDD from an AIX host system

Notes:

1. 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 37 and “Stopping the SDD server” on page 38 for more instructions. See Chapter 10, “The SDD server daemon”, on page 215 for more details about the SDD server daemon.
2. 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 Subsystem Device Driver” on page 25.

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. Type `smitty deinstall` from your desktop window to go directly to the Remove Installed Software panel.
2. Type 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 20 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.

SAN Volume Controller Preferred Node path-selection algorithm

The SAN Volume Controller is a two-controller disk subsystem. SDD distinguishes the paths to a SAN Volume Controller 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.

Dynamically changing the SDD path-selection policy algorithm

SDD 1.4.0.0 supports path-selection policies that increase the performance of a multipath-configured ESS or SAN Volume Controller 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. 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.

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. You can change the policy for an SDD device. SDD version 1.3.3.9 (or later) supports dynamic changing of the SDD devices path-selection policy.

Before changing the path-selection policy, determine the active attributes for the SDD device. Type the `lsattr -E1 vpathN` 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 issue 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 228 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. The following section shows you how to add paths to SDD devices from AIX 4.3.2 (or later) host systems.

Adding paths from AIX 4.3.2 (or later) host systems

If your host system is running AIX 4.3.2 or higher, you can issue the **addpaths** command to add paths to SDD devices of a volume group.

If your host system is running AIX 5.2.0, see “Adding more paths to the existing vpaths” on page 36.

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

If you issue the **addpaths** command to a vpath device that is in the Open state, the paths that are added are automatically in the Open state. An exception to this is when you issue the **addpaths** command to add a new path to a vpath device that has only one existing 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.

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.

Complete the following steps to add paths to SDD devices of a volume group with the **addpaths** command:

1. Issue 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/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 in one of the following ways to recognize all new **hdisk** devices. Ensure that all logical drives on the ESS or SAN Volume Controller are identified as **hdisks** before continuing.
 - Issue the **cfgmgr** command *n* times, where *n* represents the number of paths for SDD, or
 - Issue the **cfgmgr -l [scsiN/fcsN]** command for each relevant SCSI or FCP adapter.
5. Issue the **addpaths** command from the AIX command line to add more paths to the SDD devices.
6. Type 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 path” on page 226 in Chapter 11, “Using the datapath commands”, on page 217 for more information.

Adding more paths to the existing vpaths

This section applies only to package `devices.sdd.52.rte` on AIX 5.2.0.

If you are using AIX 5.2.0, there are certain limitations on the number of paths per *vpath* that you can have, depending on the number of *vpaths* you have. Use the following procedure to add more paths to configured *vpaths* on AIX 5.2.0.

1. 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 8 on page 22 for the maximum paths allowed.

For example, on a 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*.

2. Configure the extra number of **hdisks** based on step 1.

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

3. Execute **addpaths**.

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.

To perform concurrent download of LIC for the 2105 Models E10/E20 F10/F20, refer to the *2105 Model E10/E20 F10/F20 Microcode Installation Instructions* for instructions.

To perform concurrent download of LIC for the 2105 Model 800, refer to the *2105 Model 800 Microcode Installation Instructions* for instructions.

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 10, “The SDD server daemon”, on page 215 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 `sddsrv` daemon installed should apply the PTFs listed in “PTFs for APARs on AIX with Fibre Channel and the SDD server” on page 39.

Verifying if the SDD server has started

After you have installed SDD, verify if the SDD server (`sddsrv`) has automatically started by typing `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” on page 38 to proceed.

Starting the SDD server manually

If the SDD server did not start automatically after you performed the SDD installation, you can start `sddsrvc` by typing `startsrc -s sddsrvc`

Go to “Verifying if the SDD server has started” on page 37 to see if you successfully started the SDD server.

Note: The `startsrc -s sddsrvc` command starts the SDD server with 20001 as the default port number. However, the SDD server will not start on port number 20001 if another application is already using this port number. You can start the SDD server using a different port number. Go to “Changing to a different port number for the SDD server” for more information.

Changing to a different port number for the SDD server

You can change to a different port number for the SDD server.

Note: Before changing the port number, you must stop the SDD server if it is running. Go to “Stopping the SDD server” for instructions. The SDD server will not start on a port number if another application is already using the same port number. Use a port number that is not being used by other applications.

Perform the following steps:

1. Type `startsrc -a pPN -s sddsrvc` to change to a different port number. (where *PN* represents the new port number that you want to use)
For example, to start the SDD server on port 22000, type `startsrc -a p22000 -s sddsrvc`
2. Type `ps -ef | grep sddsrvc` to verify if the SDD server has automatically started on the port number that you have specified or see “Verifying if the SDD server has started” on page 37 to determine if the SDD server state is *Active*.

Tip:

- If SDD Server has started, the new port number that you specified is stored in the `/etc/sddsrvc.conf` file and will be used as default port number, until the next time that the port number is changed.

Stopping the SDD server

You can permanently disable `sddsrvc` or replace the current version of `sddsrvc` with a standalone version.

Permanently disabling the SDD server

To permanently disable `sddsrvc` on your server:

- Type `stopsrc -s sddsrvc` to stop the current version of `sddsrvc`.
- Comment out following line in system init table (`/etc/inittab`):

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

- `Sddsrvc` will not start again if the system is rebooted.

Replacing the current `sddsrvc` with a standalone version of `sddsrvc`

If you are running HACMP and are experiencing problems associated with `sddsrvc` (see “Understanding SDD support for High Availability Cluster Multi-Processing” on page 42), see Table 17 on page 48 for information about HACMP fixes that will solve the problems.

To replace the current version of sddsrv with a standalone version:

1. Type `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 37 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. Type `startsrc -a s0 -s sddsrv` to start a standalone version of sddsrv.

Starting the sddsrv daemon 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 path” on page 225 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 path” on page 225 for more information.

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

If you have Fibre Channel support and the SDD server daemon running, PTFs for the APARs shown in Table 9 must be applied to your AIX servers in order to avoid a performance degradation.

Table 9. 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)	<ul style="list-style-type: none"> • 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 the sddsrv daemon 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 35.

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 an ESS or SAN Volume Controller 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 path” on page 225 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 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 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.

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 an ESS or SAN Volume Controller, 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 SDD support for High Availability Cluster Multi-Processing

You can run the SDD in concurrent and nonconcurrent multihost environments in which more than one host is attached to the same LUNs on an ESS or SAN Volume Controller. RS/6000 and pSeries servers running High Availability Cluster Multi-Processing (HACMP) in concurrent or nonconcurrent mode are supported. Different SDD releases support different kinds of environments. (See Table 10 on page 43, Table 11 on page 44, Table 12 on page 45, Table 13 on page 46, Table 14 on page 47 and Table 15 on page 47.)

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 relationship among nodes in a cluster. This process is known as *node failover* 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.

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 SDD v1.3.3.11 (or earlier).

With SDD 1.4.0.0 (or later), there are no longer different SDD packages for HACMP running in different modes. A single package (corresponding to the AIX OS level) applies to HACMP running in different modes. For example, `devices.sdd.51.rte` supports HACMP running in both concurrent and nonconcurrent mode in the AIX 5.1 Operating System.

HACMP running in concurrent mode is supported with the `ibmSdd_432.rte` installation package for SDD 1.1.4 (SCSI only). HACMP running in concurrent mode is supported with the `ibmSdd_432.rte` installation package for SDD 1.3.1.3 (or later) (SCSI and fibre) and `ibmSdd_510.rte` installation package for SDD 1.3.1.3 (or later) (SCSI and fibre). The `ibmSdd_433.rte` installation package for SDD 1.3.1.3 (or later) and the `ibmSdd_510nchacmp.rte` installation package for SDD 1.3.1.3 (or later) are for HACMP environments only. These versions support nonconcurrent modes. However, to make the best use of the manner in which the device reserves are made, IBM recommends that you:

- Use either `ibmSdd_432.rte` installation package for an SDD level earlier than 1.4.0.0, or the `ibmSdd_510.rte` installation package for SDD 1.3.3.9 when running HACMP concurrent mode.
- Use either `ibmSdd_433.rte` installation package for an SDD level earlier than 1.4.0.0, or the `ibmSdd_510nchacmp.rte` installation package for SDD 1.3.3.9 when running HACMP in nonconcurrent mode.

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:

ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

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 10
- Table 11 on page 44
- Table 12 on page 45
- Table 13 on page 46
- Table 14 on page 47
- Table 15 on page 47

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

Table 10. Software support for HACMP in concurrent and nonconcurrent mode 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
<code>devices.sdd.43.rte</code> installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)	<ul style="list-style-type: none"> • IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • IY14682 • F model requires IY11110 	<ul style="list-style-type: none"> • IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • IY14683 • F model requires IY11480 	Not applicable

Table 10. Software support for HACMP in concurrent and nonconcurrent mode 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.51.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)	<ul style="list-style-type: none"> • IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480 	<ul style="list-style-type: none"> • IY23761 • IY26561 • (See Note.) 	<ul style="list-style-type: none"> • IY36938 • IY36933 • IY35735 • IY36951
devices.sdd.52.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)	Not applicable	Not applicable	<ul style="list-style-type: none"> • IY36938 • IY36933 • IY36782 • IY37744 • IY37746 • IY35810 • IY36951
<p>Note: For up-to-date APAR information for HACMP, go to the following Web site:</p> <p>http://service.software.ibm.com/rs6k/fixdb.html</p>			

Table 11. Software support for HACMP in concurrent mode on AIX 5.1.0 (32-bit or 64-bit kernel)

SDD version and release level	HACMP 4.4 + APARs	HACMP 4.4.1 + APARs (64-bit)	HACMP 4.5 + APARs
ibmSdd_510.rte installation package for SDD 1.3.1.3 (SCSI/FCP)	<ul style="list-style-type: none"> • IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480 	<ul style="list-style-type: none"> • IY23761 • IY26561 • (See Note.) 	<ul style="list-style-type: none"> • IY36938 • IY36933 • IY35735 • IY36951

Table 11. Software support for HACMP in concurrent mode on AIX 5.1.0 (32-bit or 64-bit kernel) (continued)

SDD version and release level	HACMP 4.4 + APARs	HACMP 4.4.1 + APARs (64-bit)	HACMP 4.5 + APARs
ibmSdd_510.rte installation package for SDD 1.3.2.9 (to SDD 1.3.3.x) (SCSI/FCP)	<ul style="list-style-type: none"> • IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480 	<ul style="list-style-type: none"> • IY23761 • IY26561 • IY25995 • (See Note.) 	<ul style="list-style-type: none"> • IY36938 • IY36933 • IY35735 • IY36951
<p>Note: For up-to-date APAR information for HACMP, go to the following Web site: http://service.software.ibm.com/rs6k/fixdb.html</p>			

Table 12. Software support for HACMP in nonconcurrent mode 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)	<ul style="list-style-type: none"> • IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480 	<ul style="list-style-type: none"> • IY23761 • IY26561 • (See Note.) 	<ul style="list-style-type: none"> • IY36938 • IY36933 • IY35735 • IY36951
ibmSdd_510nchacmp.rte installation package for SDD 1.3.2.9 (to SDD 1.3.3.x) (SCSI/FCP)	<ul style="list-style-type: none"> • IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480 	<ul style="list-style-type: none"> • IY23761 • IY26561 • IY25995 • (See Note.) 	<ul style="list-style-type: none"> • IY36938 • IY36933 • IY35735 • IY36951
<p>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: http://service.software.ibm.com/rs6k/fixdb.html</p>			

Table 13. Software support for HACMP in concurrent mode

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)	<ul style="list-style-type: none"> • IY07392 • IY03438 • IY11560 • IY08933 • IY11564 • IY12021 • IY12056 • F model requires IY11110 	<ul style="list-style-type: none"> • 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)	<ul style="list-style-type: none"> • IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • F model requires IY11110 	<ul style="list-style-type: none"> • 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)	<ul style="list-style-type: none"> • IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • F model requires IY11110 	<ul style="list-style-type: none"> • IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • F model requires IY11480 	(See Note.)
<p>Note: For up-to-date APAR information for HACMP, go to the following Web site: http://service.software.ibm.com/rs6k/fixdb.html</p>			

Table 14. Software support for HACMP in nonconcurrent mode

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)	<ul style="list-style-type: none"> • IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • IY14682 • F model requires IY11110 	<ul style="list-style-type: none"> • 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)	<ul style="list-style-type: none"> • IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • IY14682 • F model requires IY11110 	<ul style="list-style-type: none"> • IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • IY14683 • F model requires IY11480 	(See Note.)
<p>Note: For up-to-date APAR information for HACMP, go to the following Web site: http://service.software.ibm.com/rs6k/fixdb.html</p>			

Table 15. Software support for HACMP/ES on 2105 Model 800

HACMP hardware	HACMP/ES 4.4.1 + APARs (64-bit)	HACMP/ES 4.5 + APARs
2105-800 IBM 2105 TotalStorage Enterprise Storage Server (ESS) Model 800 Letter 102-201	<ul style="list-style-type: none"> • IY31627 • (See Note.) 	<ul style="list-style-type: none"> • IY32874 • (See Note.)
<p>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: http://service.software.ibm.com/rs6k/fixdb.html</p>		

Note: For the most up-to-date list of required APARs go to the following Web sites:

- www.ibm.com/storage/hardsoft/products/ess/supserver.htm
- www.service.software.ibm.com/rs6k/fixdb.html

Even though SDD supports HACMP, certain combinations of features are not supported. Table 16 on page 48 lists those combinations.

Table 16. HACMP and supported SDD features

Feature	RS/6000 or pSeries node running HACMP
ESS concurrent code load on multipath mode	Yes
SDD load balancing	Yes
SCSI	Yes
FCP (fibre)	Yes
SCSI and Fibre Channel connections to the same LUN from one host (mixed environment)	No

Notes:

1. HACMP nonconcurrent mode is supported with SDD 1.2.1.
2. SDD 1.2.2.0 requires the ESS microcode level that supports Persistent/Reserve. Do not run SDD 1.2.2.0 if your ESS microcode does not have this support.
3. HACMP nonconcurrent mode is supported with AIX 4.3.3.
4. HACMP is not supported on all models of the ESS.
5. For information about supported ESS models and required ESS microcode levels, go to the Web site www.ibm.com/storage/support

HACMP and sddsr

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 2-5 minutes intervals. It probes the idle paths at 1-minute intervals.

During the HACMP failover process, the failover 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 17. When the fixes are installed, sddsr can be re-enabled. If you have removed or commented out the line in the system /etc/inittab file that automatically starts sddsr after a system reboot, you should replace or remove the comment notation from that line in the system /etc/inittab file.

Table 17. 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 ESS 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 SAN Volume Controller or an ESS 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, type 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 ESS or SAN Volume Controller hdisks for the vpath devices. Otherwise, if the ESS or SAN Volume Controller 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 76) to convert the volume group’s physical volumes from ESS or SAN Volume Controller 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 62 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 ESS or SAN Volume Controller 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 the always the same. See “Exporting HACMP RAID concurrent-mode volume groups” on page 55

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. Type `smitty datapath_mkvg` at the command prompt.
2. A screen similar to the following displays. 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 may 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. Type `smitty importvg` at the command prompt.

A screen similar to the following displays. 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. Issue 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. Type `smitty vg` at the command prompt.
3. Select **Remove a Volume Group** from the displayed menu.

Note: A screen similar to the following displays. 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. Type `smitty datapath_extndvg` at the command prompt of one of the nodes.
3. A screen similar to the following displays. 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. Issue 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. Issue `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. Type `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 displays. 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 extending it on the current node.
7. For all the nodes sharing con_vg, do the following:
 - a. Issue **exportvg con_vg** at the command prompt.
 - b. Type **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. Type **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 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. The PTFs are listed in Table 5 on page 8.

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 58.

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. Type `smitty datapath_mkvg` at the command prompt.
2. A screen similar to the following displays. 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.

Type `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`. Type `hd2vp` to convert the volume group to `vpath` based.
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 any of the nodes in the cluster prior to attempting to retrieve the `pvid`.

- A screen similar to the following displays. 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:

- Type `smitty datapath_extendvg` at the command prompt.
- A screen similar to the following displays. Enter the information appropriate to 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:

- Type `smitty vg` at the command prompt.
- Select **Set Characteristics of a Volume Group** from the displayed menu.
- Select **Remove a Physical Volume from a Volume Group** from the displayed menu.
- A screen similar to the following displays. Enter the information appropriate to 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.
- Issue 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 35 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 higher) 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 yes for **Create Volume Group concurrent-capable?**. 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 55.

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 an ESS or to a SAN Volume Controller device.

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

Note: SDD supports secondary-system paging on ESS and SAN Volume Controller 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.

Attention: AIX does not recommend moving the primary paging space out of rootvg. Doing so may 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.

For more information about AIX paging space, go to the following Web site:

techsupport.services.ibm.com/server/support

Listing paging spaces

You can list paging spaces by typing the following:

```
lspv -a
```

Adding a paging space

You can add a paging space by typing the following:

```
mkpvs -a -n -sNN vg
```

The **mkpvs** command recognizes the following options and arguments:

- | | |
|--------------------|---|
| -a | The -a option makes the new paging space available at all system restarts. |
| -n | The -n option activates the new paging space immediately. |
| -s<i>NN</i> | The -s argument accepts the number of logical partitions (<i>NN</i>) to allocate to the new paging space. |
| <i>vg</i> | <i>vg</i> is 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, type:

```
rmpps PS01
```

Providing load-balancing and failover protection

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

Displaying the ESS or SAN Volume Controller vpath device configuration

To provide failover protection, an ESS or SAN Volume Controller vpath device must include a minimum of two paths. Both the SDD vpath device and the ESS or SAN Volume Controller hdisk devices must all 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 ESS or SAN Volume Controller 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. Type `smitty device` from your desktop window. The Devices panel displays.
2. Select **Data Path Devices** and press Enter. The Data Path Devices panel displays.
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.

Example of vpath devices with or without failover protection: vpath0, vpath1, or vpath2 has a single path and therefore does not have failover protection. The other vpath devices each have two paths and thus can provide failover

protection. The requirement for failover protection is that the vpath device, and at least two hdisk devices making it up, must be in the Available state.

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. Issue 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 ESS or SAN Volume Controller 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 77 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 220 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 60) and then add paths later by reconfiguring the ESS or SAN Volume Controller. (See “Dynamically adding paths to SDD devices of a volume group” on page 35 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. Type **smitty** at the AIX command prompt. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type **smi t** to invoke the graphical user interface.
 - 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 displays.
 3. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.

4. Select **Volume Group** and press Enter. The **Volume Groups** panel displays.
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 on page 60, 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.

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. Type `smitty` from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
 - 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 displays.
 3. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
 4. Select **Volume Groups** and press Enter. The Volume Groups panel displays.
 5. Select **Import a Volume Group** and press Enter. The Import a Volume Group panel displays.
 6. In the Import a Volume Group panel, perform the following tasks:
 - Enter the volume group you want to import.
 - Enter the one of the physical volumes that you want to import.
 - Press the **Enter** key after making the changes.

You can press the **F4** key 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 76) to convert the volume group’s physical volumes from ESS or SAN Volume Controller `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 49 for special requirements regarding unconfiguring and removing SDD `vpath` devices when running in HACMP environments.

See page 50 for an example of `hdisks`’ `pvid` 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 a nonshared 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. Type `smitty` from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
 - 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 displays.
 3. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
 4. Select **Volume Groups** and press Enter. The Volume Groups panel displays.
 5. Select **Export a Volume Group** and press Enter. The Export a Volume Group panel displays.
 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.

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 paths to a vpath pseudo device. A pseudo device loses failover protection when it has only a single path. You can use the **datapath query device** command to show the state of paths to a pseudo device. You cannot use any DEAD path for I/O operations.

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 ESS or SAN Volume Controller 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 displays:

```

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 -d1 vpath3
rmdev -d1 hdisk4
rmdev -d1 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:

- Issue **cfgmgr** once for each installed SCSI or Fibre Channel adapter.
- Issue **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 ESS or SAN Volume Controller devices is not always required. It depends on whether the ESS or SAN Volume Controller 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 ESS or SAN Volume Controller 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 ESS or SAN Volume Controller 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 use either 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 35 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 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
```

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 in 60) and then add paths later by reconfiguring the ESS or SAN Volume Controller. With a single path, failover protection is not provided. (See “Dynamically adding paths to SDD devices of a volume group” on page 35 for information about adding paths to an SDD device.)

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

1. Type **smitty** from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type **smi t** to invoke the graphical user interface.
- 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 displays.
3. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
4. Select **Volume Group** and press Enter. The **Volume Groups** panel displays.
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 Back Up a Volume Group with Data Path Devices SMIT panel” on page 75.

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 may be compromised if a file is modified during system back up. Keep system activity at a minimum during the system back-up 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 75.

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

Table 18. SDD-specific SMIT panels and how to proceed

SMIT panels	How to proceed Go to:


```

| 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 displays. 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 77 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. Type **smitty** from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type **smit** to invoke the graphical user interface.
- 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 displays.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel displays.
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 displays. 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 220 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. Type **smi tty** from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type **smi t** to invoke the graphical user interface.
- 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 displays.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel displays.
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. Type **smi tty** from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type **smi t** to invoke the graphical user interface.
- 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 displays.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel displays.
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. Type **smi tty** from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
 - 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 displays.
 3. Select **Data Path Devices** and press Enter. The Data Path Devices panel displays.
 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. Type `SMITTY` from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
 - 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 displays.
 3. Select **Data Path Devices** and press Enter. The Data Path Devices panel displays.
 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. Type `smi tty` from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
 - 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 displays.
 3. Select **Data Path Devices** and press Enter. The Data Path Devices panel displays.
 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. Type `smi tty` from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
 - 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 displays.
 3. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
 4. Select **Volume Groups** and press Enter. The **Volume Groups** panel displays.
 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. Type `smitty` from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
 - 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 displays.
 3. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
 4. Select **Volume Group** and press Enter. The Volume Group panel displays.
 5. Select **Add a Data Path Volume to a Volume Group** and press Enter.
 6. Type 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 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. Type `smitty` from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
 - 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 displays.

3. Select **Volume Groups** and press Enter. The Volume Groups panel displays.
4. Select **Set Characteristics of a Volume Group** and press Enter. The Set Characteristics of a Volume Group panel displays.
5. Select **Remove a Physical Volume from a Volume Group** and press Enter. The Remove a Physical Volume from a Volume Group panel displays.

Accessing the Back Up 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 back up a volume group with SDD devices:

1. Type `smitty` from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
- 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 displays.
3. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
4. Select **Volume Groups** and press Enter. The Volume Groups panel displays.
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 displays.
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 back up.
 - 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 back up.

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 may be compromised if a file is modified during system back up. Keep system activity at a minimum during the system back-up 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. Type `smitty` from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.

- 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 displays.
 3. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
 4. Select **Volume Groups** and press Enter. The Volume Groups panel displays.
 5. Select **Remake a Volume Group with Data Path Devices** and press Enter. The Remake a Volume Group with Data Path Devices panel displays.
 6. Enter the Restore DEVICE or FILE name that you want to restore, and press Enter. You can also use the F4 key to list all the available SDD devices, and you can select the devices or files you want to restore.

SDD utility programs

addpaths

You can use the **addpaths** command to dynamically add more paths to SDD devices while 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 path of the vpath is more than one. You can use either the Add Paths to Available Data Path Devices SMIT panel, or run the **addpaths** command from the AIX command line.

The syntax for this command is:

```
▶▶—addpaths—▶▶
```

For more information about this command, go to “Dynamically adding paths to SDD devices of a volume group” on page 35.

hd2vp and vp2hd

SDD provides two conversion scripts, **hd2vp** and **vp2hd**. The **hd2vp** script converts a volume group from ESS or SAN Volume Controller hdisks to SDD vpaths, and the **vp2hd** script converts a volume group from SDD vpaths to ESS or SAN Volume Controller hdisks. Use the **vp2hd** program when you want to configure your applications back to original ESS or SAN Volume Controller 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 ESS or SAN Volume Controller hdisk device special files before removing SDD.

The syntax for these conversion scripts is as follows:

```
▶▶—hd2vp vname—▶▶
```


▶▶—vp2hd vname—▶▶

vgname

Specifies the volume group name to be converted.

These two conversion programs require that a volume group contain either *all* original ESS or SAN Volume Controller 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 ESS or SAN Volume Controller hdisk devices can cause original ESS or SAN Volume Controller 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 hdiskN -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 rather than the vpath device. The result is that an SDD volume group partially uses SDD vpath devices, and partially uses ESS or SAN Volume Controller 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:

▶▶—dpovgfix vname—▶▶

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 30 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 vpathN0 vpathN1 vpathN2
```

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 30 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).

2105E Display all 2105 E models (ESS).

2105F Display all 2105 F models (ESS).

2105F10

Display all 2105 F10 models (ESS).

21058

Display all 2105 800 models (ESS).

2145 Display all 2145 models (SAN Volume Controller).

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 61.

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 68.

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 SDD Configuration.

The syntax is:

```
►► querysn <-d> -l device-name ◀◀
```

- l Specifies the logical number of the ESS or SAN Volume Controller 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 ESS or SAN Volume Controller 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 ESS or SAN Volume Controller 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. The `-d` flag will delete the entire contents of the `/etc/vpexclude` file, so that no devices will be excluded from the SDD configuration. 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 will 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 ESS devices in the SDD configuration” on page 17.

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 ESS or SAN Volume Controller 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 preempt-abort or clear persistent reserve service action is implemented. With preempt-abort service action not only the current persistent reserve key is preempted; it also aborts tasks on the LUN that originated from the initiators that are registered

with the preempted key. With clear service action, both persistent reservation and reservation key registrations are cleared from the device or LUN.

The following information describes in detail the syntax and examples of the **lquerypr** command.

lquerypr command

Purpose

To query and implement certain SCSI-3 persistent reserve commands on a device.

Syntax

```
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. If 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, type: `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, type: `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, type: `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, type: `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, type:

```
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 ESS or SAN Volume Controller devices directly

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

To use the SDD load-balancing and failover features and access ESS or SAN Volume Controller devices, your application must use the SDD vpath devices rather than the ESS or SAN Volume Controller hdisk devices.

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

If your application used ESS or SAN Volume Controller 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. Type `smitty` from your desktop window. The System Management Interface Tool displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
- 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 displays.

3. Select **Data Path Devices** and press Enter. The Data Path Devices panel displays.
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 type `lsvpcfg` from the command-line interface rather than using SMIT. This displays all configured SDD vpath devices and their underlying paths (hdisks).

Using ESS or SAN Volume Controller devices through AIX LVM

Attention:

- You must use the System Management Interface Tool (SMIT). The SMIT facility runs in two interfaces, nongraphical (type `smi tty` to invoke the nongraphical user interface) or graphical (type `SMIT` to invoke the graphical user interface).
- Do not use the **mkvg** command directly. Otherwise, the path failover capability could be lost.

If your application accesses ESS or SAN Volume Controller 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 ESS or SAN Volume Controller device hdisks to the SDD vpaths:

1. Determine the file systems or logical volumes that your application accesses.
2. Type `smi tty` from your desktop window. The System Management Interface Tool displays.
3. Select **System Storage Management (Physical & Logical Storage)** and press Enter. The System Storage Management (Physical & Logical Storage) panel displays.
4. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
5. Select **Logical Volume** and press Enter. The Logical Volume panel displays.
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. Type `smi tty` from your desktop window. The System Management Interface Tool displays.
- b. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
- c. Select **File Systems** and press Enter. The File Systems panel displays.
- 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 to convert the volume group from the original ESS or SAN Volume Controller hdisks to SDD vpaths:


```
hd2vp vgroupname
```
 9. When the conversion is complete, mount all file systems that you previously unmounted.

When the conversion is complete, your application now accesses ESS or SAN Volume Controller physical LUNs through SDD vpath devices. This provides load-balancing and failover protection for your application.

Migrating a non-SDD volume group to an ESS or SAN Volume Controller 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 20.
- The ESS or SAN Volume Controller 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:
 - Type `smitty` and press Enter from your desktop window. The System Management Interface Tool panel displays.

Tip:

- The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
 - 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 displays.
 - Select **Data Path Device** and press Enter. The Data Path Device panel displays.
 - 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 displays.
- 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 (ESS or SAN Volume Controller LUN) does not have a pvid written on it. Issue 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 hdiskW -a pv=clear
chdev -l vpathW -a pv=clear
```

Attention: Exercise care when clearing a pvid from a device with this command. Issuing this command to a device that *does* belong to an existing volume group can cause system failures.

You should complete the following steps to migrate a non-SDD volume group to a multipath SDD volume group in concurrent mode:

1. Add new SDD vpath devices to an existing non-SDD volume group:
 - a. Type `smitty` and press Enter from your desktop window. The System Management Interface Tool panel displays.
Tip:
 - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smi t` to invoke the graphical user interface.
 - 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 displays.
 - c. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
 - d. Select **Volume Group** and press Enter. The Volume Group panel displays.
 - e. Select **Add a Data Path Volume to a Volume Group** and press Enter.
 - f. Type 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. Mirror logical volumes from the original volume to an SDD ESS or SAN Volume Controller volume. Issue the command:

```
smitty mklvcopy
```

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. Issue the following command to synchronize all the volumes:

```
smitty syncvg
```

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. Type the following command to remove the original copy of the logical volumes from all original non-SDD physical volumes:

```
smitty rmlvcopy
```

5. Remove the original non-SDD devices from the volume group. Type the following command:

```
smitty reducevg
```

The Remove a Physical Volume panel displays. Remove all non-SDD devices.

Note: A non-SDD volume group can consist of non-ESS, non-SAN Volume Controller, ESS, or SAN Volume Controller 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 ESS or SAN Volume Controller 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. Type `smitty` and press Enter from your desktop window. The System Management Interface Tool panel displays.
Tip:
 - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can type `smit` to invoke the graphical user interface.
 - 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 displays.
 - c. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
 - d. Select **Volume Group** and press Enter. The Volume Group panel displays.
 - e. Select **Add a Data Path Volume to a Volume Group** and press Enter.
 - f. Type `vg1` in the **Volume Group Name** field. Type `vpath12` in the **Physical Volume Name** field. Press Enter.
You can also type the command:

```
extendvg4vp -f vg1 vpath12
```
2. Mirror logical volumes from the original volume to the new SDD vpath volume:
 - a. Type `smitty` and press Enter from your desktop window. The System Management Interface Tool panel displays.
 - b. Select **System Storage Management (Physical & Logical)** and press Enter. The System Storage Management (Physical & Logical) panel displays.
 - c. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
 - d. Select **Volume Group** and press Enter. The Volume Group panel displays.
 - e. Select **Mirror a Volume Group** and press Enter. The Mirror a Volume Group panel displays.
 - f. Type a volume group name. Type a physical volume name. Press Enter.
You can also type the command:

```
mirrorvg vg1 vpath12
```
3. Synchronize the logical volumes in the volume group:

- a. Type `smitty` and press Enter from your desktop window. The System Management Interface Tool panel displays.
- b. Select **System Storage Management (Physical & Logical)** and press Enter. The System Storage Management (Physical & Logical) panel displays.
- c. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
- d. Select **Volume Group** and press Enter. The Volume Group panel displays.
- e. Select **Synchronize LVM Mirrors** and press Enter. The Synchronize LVM Mirrors panel displays.
- f. Select **Synchronize by Physical Volume**.

You can also type the command:

```
syncvg -p hdisk13 vpath12
```

4. Delete copies of all logical volumes from the original physical volume:
 - a. Type `smitty` and press Enter from your desktop window. The System Management Interface Tool panel displays.
 - b. Select **Logical Volumes** and press Enter. The Logical Volumes panel displays.
 - c. Select **Set Characteristic of a Logical Volume** and press Enter. The Set Characteristic of a Logical Volume panel displays.
 - d. Select **Remove Copy from a Logical Volume** and press Enter. The Remove Copy from a Logical Volume panel displays.

You can also type the command:

```
rmlvcopy loglv01  
1 hdisk13 and rmlvcopy lv01 1 hdisk13
```

5. Remove the old physical volume from the volume group:
 - a. Type `smitty` and press Enter from your desktop window. The System Management Interface Tool panel displays.
 - b. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel displays.
 - c. Select **Volume Groups** and press Enter. The Volume Groups panel displays.
 - d. Select **Set Characteristics of a Volume Group** and press Enter. The Set Characteristics of a Volume Group panel displays.
 - e. Select **Remove a Physical Volume from a Volume Group** and press Enter. The Remove a Physical Volume from a Volume Group panel displays.

You can also type the command:

```
reducevg vg1 hdisk13
```

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. Type `'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, type:

```
trace -a -j 2F8
```

To stop the trace function, type:

```
trcstop
```

To read the report, type:

```
trcrpt | pg
```

To save the trace data to a file, type:

```
trcrpt > filename
```

Note: To perform the AIX trace function, you must have the bos.sysmgt.trace installation package installed on your system.

Error log messages

SDD logs error states into the AIX error log system. To check if SDD has generated an error log message, type the following command:

```
errpt -a | grep VPATH
```

The following list shows the SDD error log 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. Issue the **datapath** command to set the offline path to online. For more information, see Chapter 11, "Using the datapath commands", on page 217.

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.

New and modified error log messages by SDD for the HACMP environment only

The following list shows the new and modified error log messages generated by SDD installed from the `ibmSdd_433.rte` or `ibmSdd_510nchacmp.rte` installation package. The following list shows the new and modified error log messages generated by SDD installed on HACMP environments only. See “SDD persistent reserve attributes” on page 48 for more information on this release.

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.

Chapter 3. Using SDD on a Hewlett Packard host system

This chapter provides step-by-step procedures for you to install, configure, remove, and use SDD on a Hewlett-Packard (HP) host system that is attached to an ESS device or a SAN Volume Controller device. SDD 1.4.0.0 supports the coexistence of ESS devices and SAN Volume Controller devices.

For updated and additional information that is not included in this manual, please see the README file on the compact disc or go to SDD web site at:
www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

Verifying the hardware and software requirements

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

Hardware

- ESS device or a SAN Volume Controller device or both
- 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 host system go to:
storage.ibm.com/hardsoft/products/ess/supserver.htm
- 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 an ESS or a SAN Volume Controller 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
- The 2105 Model 800 does not support HP SCSI in a 32-bit mode

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 20 on page 91 or newer patches are installed on your HP-UX 11.0 host system.

Understanding how SDD works on an HP 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 ESS or SAN Volume Controller devices to your host system and attach the required SCSI or Fibre Channel adapters.

The **querysn** command can be used to exclude any ESS devices. The **querysn** command reads the unique serial number of an ESS device (sdisk) and saves the serial number in an exclude file. See “Manual exclusion of ESS devices from the SDD configuration” on page 16.

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*.

Planning for installation

Before you install SDD on your HP 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 19 on page 91 further describes the various installation scenarios and how you should proceed.

Table 19. SDD installation scenarios

Installation scenario	Description	How to proceed
Scenario 1	<ul style="list-style-type: none"> SDD is not installed. The SDD server for Expert is installed. No software application or DBMS communicates directly to sdisk interface. 	Go to: <ol style="list-style-type: none"> “Determining if the SDD 1.3.1.5 (or later) server for Expert is installed” “Installing SDD” on page 92 “Standard UNIX applications” on page 107
Scenario 2	<ul style="list-style-type: none"> SDD is not installed. The SDD server for Expert is installed. An existing application package or DBMS communicates directly to the sdisk interface. 	Go to: <ol style="list-style-type: none"> “Determining if the SDD 1.3.1.5 (or later) server for Expert is installed” “Installing SDD” on page 92 “Using applications with SDD” on page 107
Scenario 3	<ul style="list-style-type: none"> SDD is installed. The SDD server for Expert is installed. 	Go to <ol style="list-style-type: none"> “Determining if the SDD 1.3.1.5 (or later) server for Expert is installed” “Upgrading SDD” on page 96

For SDD to operate properly, ensure that the patches shown in Table 20 are installed on your HP-UX 11.0 host system.

Note: Do *not* install the patches noted in Table 20 on HP-UX 11i (Fibre Channel) host systems.

Table 20. Patches necessary for proper operation of SDD on HP-UX 11.0

Application mode	HP 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 IO Subsystem Cumulative patch
32-bit	PHKL_23001	Signal, threads, spinlock, scheduler, IDS, q3p
32-bit	PHKL_23406	Probe, sysproc, shm, 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
64-bit	PHKL_21381	Fibre Channel Mass Storage Driver

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 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, type:

```
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 `SDDsrvHPbb_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 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:
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/swexpert

For more information about the SDD server daemon, go to “SDD server daemon” on page 102.

Installing SDD

Before you install SDD, make sure that you have root access to your HP host system and that all the required hardware and software is ready.

Perform the following steps to install SDD on your HP 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 3 or Figure 4.

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

Figure 3. IBMsdd Driver 32-bit

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

Figure 4. 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 1.4.0.0 is changed from a static driver to a dynamic loadable kernel module (DLKM) driver. The system will not restart after SDD is installed.

SDD 1.4.0.0 (or later) does not support the HP 32-bit 11i PA1.1 system because of a known problem with the HP DLKM module on the 32bit 11i PA1.1 HP-UX system. Currently, only static-linked SDD 1.3.2.8 (or earlier) supports 32bit 11i PA1.1 system.

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.

After you have verified that SDD has been installed, vary on the volume groups with the following command:

```
vgchange -a y /dev/vg
```

where *vg* is the name of the volume group.

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 21. SDD components installed for HP 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
sddsrv.conf	/etc/	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

Table 22. System files updated for HP host systems

File	Location	Description
vpath	/usr/conf/master.d	Master configuration file
vpath	/stand/system.d	System configuration file

Table 23. SDD commands and their descriptions for HP host systems

Command	Description
cfgvpath [-c]	Configures vpath 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.
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 ESS sdisks into SDD vpaths.
vp2hd	Converts a volume group from SDD vpaths into ESS sdisks.
rmvpath [-all, -vpathname]	Removes vpath devices from the configuration.
gettrace	Debug tool, gets trace information when problem happens.

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 107 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 107 and read the information specific to the application 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/sddsrv.conf

During installation, /opt/IBMsdd/bin/cfgvpath program is initiated to create vpaths in the /dev/dsk and /dev/rdisk 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/rdisk and /dev/dsk. The device is named according to the SDD number. A device with a number of 0 would be /dev/rdisk/vpath1.

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 97 and then go to “Installing SDD” on page 92.

Upgrading SDD 1.3.0.2 (or earlier) to SDD 1.4.0.0 (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 1.4.0.0 (or later):

1. Uninstall SDD 1.3.0.2 (or earlier). Since 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 1.4.0.0 (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`
 - 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 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.

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. Unmount all SDD file systems.
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, vary off the volume groups using the following command:

```
vgchange -a n /dev/vg
```

where *vg* is the volume group name.

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 92.

Configuring SDD

This section provides information necessary to configure SDD.

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 typing:

```
/opt/IBMsdd/bin/cfgvpath -c
```

2. Restart the system by typing:

```
shutdown -r 0
```

Converting a volume group

SDD provides two conversion scripts:

hd2vp The **hd2vp** script converts a volume group from ESS or SAN Volume Controller sdisks into SDD vpaths.

The syntax for **hd2vp** script is as follows:

```
hd2vp vname
```

vp2hd The **vp2hd** script converts a volume group from SDD vpaths into ESS or SAN Volume Controller sdisks. Use the **vp2hd** program when you want to configure your applications back to original ESS or SAN Volume Controller sdisks.

The syntax for **vp2hd** script is as follows:

```
vp2hd vname
```

These two conversion programs require that a volume group contain either *all* original ESS or SAN Volume Controller 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 supports path-selection policies that increase the performance of multipath-configured ESS devices and SAN Volume Controller 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.

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. Type **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 228 for more information about the **datapath set device policy** command.

SAN Volume Controller Preferred Node path-selection algorithm

The SAN Volume Controller is a two-controller disk subsystem. SDD distinguishes the paths to a SAN Volume Controller 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

For SDD 1.4.0.0 (or later), the output of some of the datapath commands has changed. See Chapter 11, "Using the datapath commands", on page 217 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

To support both ESS and the SAN Volume Controller, the following change is made in SDD for the **datapath query device** command. The locations of Serial and Policy are swapped because the SAN Volume Controller serial is too long to fit in the first line.

Example output from **datapath query device** command issued in SDD 1.3.3.11 (or earlier):

```
Dev#: 3      Device Name: vpath5  Type: 2105800      Serial: 14123922
Policy: Optimized
=====
Path#      Adapter H/W Path      Hard Disk      State   Mode   Select  Error
0          0/7/0/0      c19t8d1      OPEN  NORMAL  3869815  0
1          0/7/0/0      c13t8d1      OPEN  NORMAL  3872306  0
2          0/3/0/0      c17t8d1      OPEN  NORMAL  3874461  0
3          0/3/0/0      c11t8d1      OPEN  NORMAL  3872868  0
```

Example output from **datapath query device** command issued in SDD 1.4.0.0 (or later): (This example shows a SAN Volume Controller 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 10, “The SDD server daemon”, on page 215 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 typing `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 92 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 typing `nohup /sbin/sddsrv&`

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

Note: The `nohup /sbin/sddsrv&` command starts the SDD server with 20001 as the default port number. However, the SDD server will not start on port number 20001 if another application is already using this port number. You can start the SDD server using a different port number. Go to “Changing to a different port number for the SDD server” for more information.

Changing to a different port number for the SDD server

You can change to a different port number for the SDD server.

Note: You must stop the SDD server if it is running before changing the port number. Go to “Stopping the SDD server” on page 103 for instructions.

Perform the following steps:

1. Type `nohup /sbin/sddsrv pPN &` to change to a different port number. (where *PN* represents the new port number that you want to use; the `&` argument starts the SDD server in the background.)
For example, to start the SDD server on port 22000, type `nohup /sbin/sddsrv p22000 &`
2. Go to “Verifying if the SDD server has started” to verify if the SDD server has automatically started on the port number that you have specified.

Tip:

- The SDD server will not start on a port number if another application is already using the same port number. Use a port number that is not being used by other applications.
- If SDD Server has started, the new port number that you specified is stored in the `/etc/sddsrv.conf` file and will be used as the default port number, until the next time that the port number is changed.

Stopping the SDD server

Perform the following steps to stop the SDD server:

1. Query the process for SDD server by typing 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 typing 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 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 98.

Understanding the SDD error recovery policy

The new SDD error recovery policy (1.3.2.2 or higher) differs from the prior 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 path” on page 226 for more information.

How to import and export volume groups

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” on page 106.
- Create the volume group device directory. See “Creating the volume group device directory” on page 106.
- Create the group special file. See “Creating the group special file” on page 106.

For more information about the **vgimport** command, see “Importing volume groups” on page 106.

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, type:

```
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* file to node 2, type:

```
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, type:

```
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, type:

```
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 (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, type:

```
vgimport -p -v -s -m /tmp/vgpath1.map vgpath1
```

where `/tmp/vgpath1.map` represents your mapfile, and `vgpath1` represents the path name of volume group that you want to 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 file server
- Oracle

Standard UNIX applications

If you have not already done so, install SDD using the procedure in “Installing SDD” on page 92. When this is done, SDD resides above the HP 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.

Type the following command to determine the major number:

```
# lsdev | grep lv
```

A message similar to the following is displayed:

```
64          64          lv          1vm
```

The first number in the message is the major number of the character device, which is the number that you want to use.

2. Create a device node for the logical volume device.

Note: If you do not have any other logical volume devices, you can use a minor number of `0x010000`. In this example, assume that you have no other logical volume devices. A message similar to the following is displayed:

```
# mknod group c 64 0x010000
```

Create a physical volume by performing the procedure in step 3 on page 109.

- a. Create a subdirectory in the `/dev` directory for the volume group.

Type 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.

Type the following command to change to the `/dev` directory:

```
# cd /dev/vgibm
```

- c. Create a device node for the logical volume device.

Type the following command to re-create the physical volume:

```
# pvcreate /dev/rdisk/vpath1
```

A message similar to the following is displayed:


```
Physical volume "/dev/rdisk/vpath1" has been successfully created.
```

In this example, the SDD device associated with the underlying disk is vpath1. Verify the underlying disk by typing 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.

Type the following command to create a physical volume:

```
# pvcreate /dev/rdisk/vpath1
```

4. Create a volume group.

Type the following command to create a volume group:

```
# vgcreate /dev/vgibm /dev/dsk/vpath1
```

5. Create a logical volume.

Type the following command to create logical volume lv01:

```
# lvcreate -L 100 -n lv01 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 filesystem, type the following command to create a file system on the volume group:

```
# newfs -F HFS /dev/vgibm/rlv01
```

b. If you are using a VXFS filesystem, type the following command to create a file system on the volume group:

```
# newfs -F VXFS /dev/vgibm/rlv01
```

c. Mount the logical volume.

This process assumes that you have a mount point called `/mnt`.

7. Mount the logical volume.

Type the following command to mount the logical volume lv01:

```
# mount /dev/vgibm/lv01 /mnt
```

Attention: In some cases it may 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:

docs.hp.com

and click on *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, type the following command to unmount logical volume `lv01`:

```
# umount /dev/vgibm/lv01
```

Next, remove the logical volume.

For example, type the following command to remove logical volume `lv01`:

```
# lvremove /dev/vgibm/lv01
```

A message similar to the following is displayed:

```
The logical volume "/dev/vgibm/lv01" is not empty;
do you really want to delete the logical volume (y/n)
```

Type `y` and press Enter. A message similar to the following is displayed:

```
Logical volume "/dev/vgibm/lv01" has been successfully removed.
Volume Group configuration for /dev/vgibm has been saved in
/etc/lvmconf/vgibm.conf
```

When prompted to delete the logical volume, type `y`.

2. Remove the existing volume group.

Type 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 `lv01` under a volume group `vgibm`, which is currently using the disk directly, (for example, through path `/dev/path /dev/dsk/c3t4d0`). You want to convert logical volume `lv01` to use SDD.

1. Determine the size of the logical volume.

Type the following command to determine the size of the logical volume:

```
# lvdisplay /dev/vgibm/lv01 | 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.

Type the following command to re-create the physical volume:

```
# pvcreate /dev/rdisk/vpath1
```

A message similar to the following is displayed:

```
Physical volume "/dev/rdisk/vpath1" has been successfully created.
```

In this example, the SDD device associated with the underlying disk is vpath1. Verify the underlying disk by typing 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.

Type 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

Type 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.

Attention: The timeout values for the logical volume manager must be set correctly 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. Type `lvdisplay /dev/vgibm/lvol1` and press Enter. If the timeout value is not default, type `lvchange -t 0 /dev/vgibm/lvol1` 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. Type `pvchange -t 240 /dev/dsk/vpathn` and press Enter. (*n* refers to the vpath number.) If you are not sure about the vpath number, type `/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.

Attention: In some cases it may 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`, `pvchange`, or `vgreduce`, see the following web site:

docs.hp.com

and click on *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 92.
2. Determine which SDD (`vpathN`) volumes 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, type the following command:

```
# newfs /dev/rdisk/vpathN
```

In this example, *N* is the SDD device instance of the selected volume. Create mount points for the new file systems.

4. Install the file systems into the directory `/etc/fstab`. 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 92.
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 92.
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/rdisk/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/rdisk/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 92.
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/rdisk/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/rdisk/c1t4d0).
3. Launch the sqlplus utility.
4. Type the command:

```
select * from sys.dba_data_files;
```

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, which is for planning purposes:

Oracle Device Link	File Attributes			SDD Device Link
	Owner	Group	Permissions	
/dev/rdisk/c1t1d0	oracle	dba	644	/dev/rdisk/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 92.
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 typing the following command:

```
/opt/IBMsdd/bin/showvpath
```
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 4. Using SDD on a Linux host system

This chapter provides step-by-step procedures on how to install, configure, use, and remove SDD on 32-bit, Intel-based Red Hat and SuSE Linux host systems that are attached to an ESS or a SAN Volume Controller. For updated and additional information that is not included in this chapter, see the README file on the compact disc or visit the SDD Web site:
www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

Verifying the hardware and software requirements

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

Hardware

- ESS device or a SAN Volume Controller or both
- 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.storage.ibm.com/hardsoft/products/ess/supserver.htm

- 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

- 32-bit, Intel-based, Red Hat 7.x Linux operating system running:
 - Red Hat 7.x kernel-2.4.18-24.7.xbigmem kernel and kernel-2.4.18-27.7.xbigmem
- Red Hat Linux Advanced Server 2.1 with:
 - 2.4.9-e.10enterprise
 - 2.4.9-e.12enterprise
 - 2.4.9-e.16enterprise
- Red Hat Linux Advanced Server 2.1 with the 2.4.9-e.16summit kernel
- SuSE SLES7 with k_smp-2.4.18-224 kernel rpm
- SuSE SLES8 with kernel k_smp-2.4.19-233 in i686

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)
- System start from an SDD pseudo device
- EXT3 filesystem on an SDD pseudo device
- System paging file on an SDD pseudo device

- Single-path mode during concurrent download and installation of 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 ESS 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

Before you install SDD, configure your SAN Volume Controller 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 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*.

Refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Host Systems Attachment Guide* for information on working around Linux LUN limitations.

Configuring Fibre Channel adapters for ESS

You must configure the Fibre Channel adapters and the adapters' 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.

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. Type `mount /dev/cdrom /mnt/cdrom` to mount the CD-ROM drive.
4. Type `cd /mnt/cdrom` to change to the `/mnt/cdrom` directory.
5. If you're running Red Hat, type `cd redhat`; if you're running SuSE, type `cd suse`, and then type `ls` to display the name of the package.
6. Type `rpm -qpi 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. Type `rpm -iv IBMsdd-N.N.N.N-1.i686.redhat.rpm` or type `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. Type `mount /dev/cdrom /mnt/cdrom` to mount the CD-ROM drive.
4. Type `cd /mnt/cdrom` to change to the `/mnt/cdrom` directory.
5. If you're running Red Hat, type `cd redhat`; if you're running SuSE, type `cd suse`, and then type `ls` to display the name of the package.
6. Type `rpm -qpi 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. Type `rpm -U IBMsdd-N.N.N.N-1.i686.redhat.rpm` or type `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 24 lists the SDD driver and its major component files.

Table 24. SDD components for a Linux host system

File name	Location	Description
sdd-mod.o-xxx	/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
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 ESS or SAN Volume Controller to which your host system is attached and that the ESS or SAN Volume Controller 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 25.

Table 25. 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/bin` directory to go to the directory where you have installed SDD utility files.
3. Change to the `/opt/IBMsdd` directory and type `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.

4. Type `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)
```

5. Type `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.
```

6. Type `cat /proc/scsi/scsi` to verify that all expected devices (LUNs) were detected and configured.

Type `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,
:
:

```

7. Type `cfgvpath query` to verify that you have configured the SCSI disk devices that you allocated and configured for SDD.

After you issue the **cfgvpath query** command, a message similar to the following is displayed. This example output is for a system with ESS and SAN Volume Controller LUNs.

```

/dev/sda ( 8, 0) host=0 ch=0 id=0 lun=0 vid=IBM pid=DDYS-T36950M serial=xxxxxxxxxxx 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/sdj ( 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=1 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=1 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
:

```

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.

8. Type `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.

Type `cfgvpath ?` or `rmvpath ?` for more information about the `cfgvpath` or `rmvpath` command.

9. Verify the vpath configuration by typing `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 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/sdab /dev/sdbe
004 vpathe ( 247, 64) 600507680183000a800000000000000a = /dev/sdj /dev/sdt /dev/sdad /dev/sdan /dev/sdbf /dev/sdba /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/sdch /dev/sdcq
012 vpathm ( 247, 192) 600507680183000a8000000000000012 = /dev/sdr /dev/sdb /dev/sdai /dev/sdav /dev/sdbn /dev/sdbx /dev/sdci /dev/sdcr
013 vpathn ( 247, 208) 600507680183000a8000000000000013 = /dev/sds /dev/sdac /dev/sdam /dev/sdaw /dev/sdbo /dev/sbvy /dev/sdci /dev/sdcs

```


See Chapter 11, “Using the datapath commands”, on page 217 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 121.

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 SLES 8 system startup:

1. Log on to your Linux host system as the root user.
2. Type `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. Type `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 typing:

```
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. Type `insserv sdd`
3. Restart your host system so that SDD is loaded and configured.

If necessary, you can disable the startup option by typing:

```
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 220 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 v1.4.0.0 supports path-selection policies that increase the performance of multipath-configured ESS devices and SAN Volume Controller 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.

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. Type **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 228 for more information about the **datapath set device policy** command.

Dynamically adding paths to SDD devices of a volume group

You can issue the **addpaths** command to add paths to SDD devices of a volume group. For SDD to discover new paths for a volume group, 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. **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. Type `lsmod` to list the modules and version of SDD that is loaded.
3. Type `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. Type `rpm -e IBMsdd` to remove the SDD package.
5. Verify the SDD removal by typing either `rpm -q IBMsdd` or `rpm -qI IBMsdd`.

If you successfully removed SDD, output similar to the following is displayed:

```
package IBMsdd is not installed
```

Note: The `rmmmod` 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 10, “The SDD server daemon”, on page 215 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 typing `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. Issue the command `telinit q`.
4. Follow the directions in “Verifying if the SDD server has started” to confirm that the SDD server started successfully.

Note: This will start the SDD server with 20001 as the default port number. However, the SDD server will not start on port number 20001 if another application is already using this port number. You can start the SDD server using a different port number. Go to “Changing to a different port number for the SDD server” for more information.

Changing to a different port number for the SDD server

You can change to a different port number for the SDD server.

Note: The SDD server will not start on a port number if another application is already using the same port number. Use a port number that is not being used by other applications.

Perform the following steps:

1. Edit `/etc/inittab` and change the `IBMsdd` entry to use the selected port number.

For example: (The following example uses port 22000.)
`srv:345:respawn:/opt/IBMsdd/bin/sddsrv p22000 > /dev/null 2>&1`

2. Save the file.
3. Run **telinit q**.
4. Run **killall sddsrv**.

Tip: If SDD Server has started, the new port number that you specified is stored in the `/etc/sddsrv.conf` file and will be used as default port number until the next time that the port number is changed.

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. Run **telinit q**.

See “Verifying if the SDD server has started” on page 128 to verify that the SDD server is not running. If `sddsrv` is not running, no processes will be listed when you type `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 ESS device or SAN Volume Controller devices 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 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 3 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 path” on page 225 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.

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 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/vpathX` 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/vpathXY` are also created for each partition of the multipath device (where `Y` represents the corresponding partition number). When a file system or user application wants to use the logical device, it should refer to `/dev/vpathXY` (for example, `/dev/vpatha1` or `/dev/vpathbc7`) as its multipath logical device. All I/O management, statistics, and failover processes of the logical device follow those of the whole device.

The following output demonstrates how the partitions are named:

```
brw-r--r-- 1 root  root  247,  0 Apr 2 16:57 /dev/vpatha
brw-r--r-- 1 root  root  247,  1 Apr 2 16:57 /dev/vpatha1
brw-r--r-- 1 root  root  247, 10 Apr 2 16:57 /dev/vpatha10
brw-r--r-- 1 root  root  247, 11 Apr 2 16:57 /dev/vpatha11
brw-r--r-- 1 root  root  247, 12 Apr 2 16:57 /dev/vpatha12
brw-r--r-- 1 root  root  247, 13 Apr 2 16:57 /dev/vpatha13
brw-r--r-- 1 root  root  247, 14 Apr 2 16:57 /dev/vpatha14
brw-r--r-- 1 root  root  247, 15 Apr 2 16:57 /dev/vpatha15
brw-r--r-- 1 root  root  247,  2 Apr 2 16:57 /dev/vpatha2
brw-r--r-- 1 root  root  247,  3 Apr 2 16:57 /dev/vpatha3
brw-r--r-- 1 root  root  247,  4 Apr 2 16:57 /dev/vpatha4
brw-r--r-- 1 root  root  247,  5 Apr 2 16:57 /dev/vpatha5
brw-r--r-- 1 root  root  247,  6 Apr 2 16:57 /dev/vpatha6
brw-r--r-- 1 root  root  247,  7 Apr 2 16:57 /dev/vpatha7
brw-r--r-- 1 root  root  247,  8 Apr 2 16:57 /dev/vpatha8
brw-r--r-- 1 root  root  247,  9 Apr 2 16:57 /dev/vpatha9
```

Note:

- SDD does not support the use of root (`/`), `/var`, `/usr`, `/opt`, `/tmp`, and swap partitions.
- 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 IO 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 5. Using SDD on a Novell host system

Attention: SDD does not support Novell host systems attached to a SAN Volume Controller.

This chapter provides step-by-step procedures on how to install, configure, upgrade, and remove SDD on a Novell host system (Novell 5.1 or Novell 6.0) that is attached to an ESS. The SDD for Novell 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 compact disc or visit the SDD Web site at:
www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

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/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
- Novell Netware 5.1 SP4 running on the server or Novell Netware 6 SP1 or SP2 running on the server
- 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

Supported environments

SDD supports:

- Novell Netware 5.1 SP4
- Novell Netware 6 SP1 or SP2
- Novell Cluster Services 1.01 for Novell Netware 5.1 is supported on Fibre Channel and SCSI devices.

- Novell Cluster Services 1.6 for Novell Netware 6.0 is supported only for Fibre Channel devices.

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 Novell host system, go to the following Web site:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

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/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 Novell 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 Novell host system go to the following Web site:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

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 Novell 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
```

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 Novell host system.

Refer to the *IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide* for working around Novell 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 Novell host system.

Refer to the *IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide* for information about working around Novell LUN limitations.

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. Type `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.

Configuring SDD

To load the SDD module, type the command `load SDD`.

To unload the SDD module, type the command `unload SDD`.

Displaying the current version of the SDD

Use the `modules SDD` command 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 11, "Using the datapath commands", on page 217 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 those

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 11, "Using the datapath commands", on page 217

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).

Module ID

Novell has assigned SDD.CDM module ID 0x7B0.

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 Novell Netware 6.0

To set up clustering in Novell Netware 6.0, follow the steps described in the Novell Cluster Services document available online at:
www.novell.com/documentation/lg/ncs6p/index.html

Chapter 6. Using SDD on a Sun host system

This chapter provides step-by-step procedures on how to install, configure, remove, and use SDD on a Sun host system that is attached to an ESS device or to a SAN Volume Controller. For updated and additional information not included in this manual, see the README file on the CD or visit the SDD Web site:

www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

Verifying the hardware and software requirements

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

Hardware

- ESS device or SAN Volume Controller or both
- 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 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.

SDD requires enabling the host adapter dynamic binding feature to have the same system device names for the same LUNs.

For information on the SCSI or Fibre Channel adapters that you can use on your Sun host system, go to the following Web site:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Software

SDD supports:

- ESS on a SPARC system running Solaris 2.6, Solaris 7, or Solaris 8.
- SAN Volume Controller on a SPARC system running 64 bit Solaris 8.

Supported environments

SDD supports the following environments:

- 32-bit applications on Solaris 2.6
- 32-bit and 64-bit applications on Solaris 7 and Solaris 8

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 Sun host system

SDD resides above the Sun 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 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

Before you install SDD, configure your SAN Volume Controller 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*.

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 Sun 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, type:

```
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 SDDsrvSUNbb_yymmdd.pkg. In this version, *bb* represents 32 or 64 bit, and *yymmdd* represents the date of the installation package. For ESS Expert V2R1, the 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 Sun 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:
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/swexpert

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

Planning for installation

Before you install SDD on your Sun 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 (DBMS)

You can install SDD in three different ways. The way that you choose depends on the kind of software that you have installed:

Table 26 further describes the various installation scenarios and how you should proceed.

Table 26. SDD installation scenarios

Installation scenario	Description	How to proceed
Scenario 1	<ul style="list-style-type: none"> • SDD is not installed. • No volume managers are installed. • No software application or DBMS is installed that communicates directly to the <i>sd</i> interface. 	Go to: <ol style="list-style-type: none"> 1. “Installing SDD” on page 144 2. “Standard UNIX applications” on page 157

Table 26. SDD installation scenarios (continued)

Scenario 2	<ul style="list-style-type: none"> SDD is not installed. An existing volume manager, software application, or DBMS is installed that communicates directly to the sd interface. 	Go to: <ol style="list-style-type: none"> “Installing SDD” “Using applications with SDD” on page 157
Scenario 3	SDD is installed.	Go to “Upgrading the SDD” on page 150.

Table 27 lists the installation package file names that come with SDD.

Table 27. Operating systems and SDD package file names

Operating System	Package file name
Solaris 2.6, Solaris 7, Solaris 8	sun32bit/IBMsdd
Solaris 7	sun64bit/IBMsdd
Solaris 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 Sun 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 Sun host system:

Note: If the OS is Solaris 2.6, the OS bit-level is 32bit.

If the OS is Solaris 7 or Solaris 8, 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. Type **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/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/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. Type 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, type 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.

Attention: Do not restart between the uninstallation and the reinstallation of SDD.

Perform the following steps to uninstall SDD:

1. Unmount all SDD file systems.
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. Type # 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. Type 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. Type 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/sddsrv.conf
/etc/rcS.d/S20vpath-config
/etc/defvpath
## Updating system information.

Removal of <IBMsdd> was successful.

```

Attention: Do not restart at this time.

After you successfully uninstall SDD, the first part of the procedure for upgrading the SDD is complete. To complete the upgrade, you now need to reinstall SDD. See “Installing SDD” on page 144 for detailed procedures.

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. Type the command:

```
# shutdown -i6 -g0 -y
```

SDD devices are found in the /dev/rdisk 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/rdisk/vpath0a where *a* denotes the slice. Therefore, /dev/rdisk/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 28. SDD components installed for Sun 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
S20vpath-config	/etc/rcS.d	Boot initialization script Note: This script must come before other LVM initialization scripts

Table 28. SDD components installed for Sun host systems (continued)

sddsrv	/sbin/sddsrv	SDD server daemon
sddsrv.conf	/etc/sddsrv.conf	SDD server config file

Table 29. System files updated for Sun 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 30. SDD commands and their descriptions for Sun host systems

Command	Description
cfgvpath [-c]	<p>Configures vpath devices using the following process:</p> <ol style="list-style-type: none"> 1. Scan the host system to find all devices (LUNs) that are accessible by the SUN 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. <p>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.</p> <p>Note: cfgvpath -c updates the configuration file but does not update the kernel. To update the kernel, you need to reboot.</p> <p>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*</p> <p>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.</p>
cfgvpath -r	Reconfigures vpath devices if vpath devices exist. See "Option 2: Dynamic reconfiguration" on page 151. If no vpath devices exist, use cfgvpath without -r option.
showvpath	Lists all SDD devices and their underlying disks

Table 30. SDD commands and their descriptions for Sun host systems (continued)

Command	Description
vpathmkdev	<p>Creates files vpathMsN in the /dev/dsk/ and /dev/rdisk/ directories by creating links to the pseudo-vpath devices /devices/pseudo/vpathdd* that are created by the SDD driver.</p> <p>Files vpathMsN in the /dev/dsk/ and /dev/rdisk/ directories provide block and character access to an application the same way as the cxydzsn devices created by the system.</p> <p>vpathmkdev is executed automatically during SDD package installation and should be executed manually to update files vpathMsN after hardware reconfiguration.</p>
datapath	SDD driver console command tool
rmvpath [all, vpathname]	Removes vpath devices from the configuration. See “Option 2: Dynamic reconfiguration” on page 151.

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 157 and read the information specific to the application that you are using.

Upgrading the SDD

Upgrading SDD consists of uninstalling and reinstalling the IBMsdd package. If you are upgrading SDD, go to “Uninstalling SDD” on page 147 and then go to “Installing SDD” on page 144.

Configuring SDD

Before you start the SDD configuration process, make sure that you have successfully configured the ESS or SAN Volume Controller to which your host system is attached and that the ESS or SAN Volume Controller 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, type `shutdown -i0 -g0 -y` and press Enter. If you do not have a console attached to your host, type `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 typing `boot -r` and pressing Enter at the OK prompt.
3. Run the SDD configuration utility in the directory `/opt/IBMsdd/bin` to reconfigure SDD.
4. Shut down the system. Type `shutdown -i6 -g0 -y` and press Enter.
5. After the restart, change to the `/opt/IBMsdd/bin` directory by typing:


```
cd /opt/IBMsdd/bin
```
6. For Solaris 7 and Solaris 8:
 - a. Type `devfsadm` and press Enter to reconfigure all the drives.

For Solaris 6:

 - a. Type `drvconfig` and press Enter.
 - b. Type `devlinks` and press Enter to reconfigure all the drives.
7. Type `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`

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 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.

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 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.

- 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 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.

Dynamically changing the SDD path-selection policy algorithm

SDD v1.4.0.0 supports path-selection policies that increase the performance of a multipath-configured ESS 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. 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.

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 228 for more information about the **datapath set device policy** command.

SAN Volume Controller Preferred Node path-selection algorithm

The SAN Volume Controller is a two-controller disk subsystem. SDD distinguishes the paths to a SAN Volume Controller 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 Sun 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's 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 Sun 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.

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 *multiple-path 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 path” on page 225 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 *multiple-path 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 10, “The SDD server daemon”, on page 215 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 typing `ps -ef | grep sddsrv`

If the SDD server (`sddsrv`) has automatically started, the output will display the process number on which `sddsrv` has started.

If the SDD Server has *not* started, go to “Starting the SDD server manually” on page 156.

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 sddsr, use the following process to start sddsr:

1. Edit `/etc/inittab` and verify the sddsr entry. It should be like one of the following lines:

```
srv:234:respawn:/sbin/sddsr > /dev/null 2>&1
```

or

```
srv:234:respawn:/sbin/sddsr pPN > /dev/null 2>&1
```

where PN is the port number on which sddsr should be started.

2. Save the file `/etc/inittab`.
3. Execute **init q**.
4. Follow the directions in “Verifying if the SDD server has started” on page 155 to confirm that the SDD server started successfully.

Note: This will start the SDD server on port number 20001 if no port number is specified or start the SDD server on the specified port number if PN is supplied. However, the SDD server will not start if another application is already using the same port number. You can start the SDD server using a different port number. Go to “Changing to a different port number for the SDD server” for more information.

Changing to a different port number for the SDD server

You can change to a different port number for the SDD server.

Note: The SDD server will not start on a port number if another application is already using the same port number. Use a port number that is not being used by other applications.

Perform the following steps:

1. Edit `/etc/inittab` and change the sddsr entry to use the selected port number. The following example uses port 22000.

```
srv:234:respawn:/sbin/sddsr p22000 > /dev/null 2>&1
```

2. Save the file.
3. Execute **init q**.
4. Execute **kill -9** pid of sddsr.

Note: When the SDD Server has started, the new port number that you specified is stored in the `/etc/sddsr.conf` file and will be used as default port number until the next time that the port number is changed.

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/sddsr > /dev/null 2>&1
```

2. Save the file.
3. Execute **init q**.
4. Check if sddsr is running by executing **ps -ef |grep sddsr**. If sddsr is still running, execute **kill -9** pid of sddsr.

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 that these software applications or DBMSs access 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
- Network File Systems
- Oracle

Standard UNIX applications

If you have not already done so, install SDD using the procedure in “Installing SDD” on page 144. 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 144 section.
2. Determine which SDD (`vpathN`) volumes 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, type the following command:

```
# newfs /dev/rdisk/vpathNs
```

In this example, *N* is the SDD device instance of the selected volume. Create mount points for the new file systems.

5. Install the file systems into the `/etc/fstab` directory. Click yes in the **mount at boot** field.
6. Install the file system mount points into the directory `/etc/exports` for export.
7. Restart the system.

Installing SDD on a system that already has the Network File System file server

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.

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 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 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 144 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 user’s `ORACLE_BASE` and `ORACLE_HOME` environment variables to be directories of this file system.
 - c. Create two more SDD-resident file systems on two other SDD volumes. Each of the resulting three mount points should have a subdirectory named `oradata`. The subdirectory is used as a control file and redo log location for the installer’s default database (a sample database) as described in the *Installation Guide*. Oracle recommends using raw partitions for redo logs. To use SDD raw partitions as redo logs, create symbolic links from the three redo log locations to SDD raw device links that point to the slice. These files are named `/dev/rdisk/vpathNs`, where *N* is the SDD instance number, and *s* is the partition ID.
3. Determine which SDD (`vpathN`) volumes you will use as Oracle8 database devices.
4. Partition the selected volumes using the Solaris format utility. If Oracle8 is to use SDD raw partitions as database devices, be sure to leave sector 0/disk cylinder 0 of the associated volume unused. This protects UNIX disk labels from corruption by Oracle8.
5. Ensure that the Oracle software owner has read and write privileges to the selected SDD raw partition device files under the `/devices/pseudo` directory.
6. Set up symbolic links in the `oradata` directory under the first of the three mount points. See step 2 on page 159. Link the database files to SDD raw device links (files named `/dev/rdisk/vpathNs`) that point to partitions of the appropriate size.
7. Install the Oracle8 server following the instructions in the *Oracle Installation Guide*. Be sure to be logged in as the Oracle software owner when you run the `oraInst /m` command. Select the **Install New Product - Create Database Objects** option. Select **Raw Devices** for the storage type. Specify the raw device links set up in step 2 for the redo logs. Specify the raw device links set up in step 3 for the database files of the default database.
8. To set up other Oracle8 databases, you must set up control files, redo logs, and database files following the guidelines in the *Oracle8 Administrator’s Reference*. Make sure any raw devices and file systems that you set up reside on SDD volumes.
9. Launch the `sqlplus` utility.
10. Issue the **create database** SQL command, specifying the control, log, and system data files that you have set up.

11. Issue the **create tablespace** SQL command to set up each of the temp, rbs, tools, and users database files that you created.
12. Issue the **create rollback segment** SQL command to create the three redo log files that you set. For the syntax of these three **create** commands, see the *Oracle8 Server SQL Language Reference Manual*.

Installing an SDD on a system that already has Oracle in place

The installation procedure for a new SDD installation differs depending on whether you are using a file system or raw partitions for your Oracle database.

If using a file system: 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 144 section.
3. Change to the directory where you installed the SDD utilities. Issue the **showvpath** command.
4. Check the display to see whether you 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/ct8d2s2`. (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:

```
vpath0c
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 Sun identifiers:

```
mount /dev/dsk/c1t3d2s4 /oracle/mp1
```

you now use the following SDD partition identifiers:

```
mount /dev/dsk/vpath2e /oracle/mp1
```

For example, assume you found that `vpath2c` was 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/rdisk/cntndn` files).

All Oracle8 control, log, and data files are accessed either directly from mounted file systems or through links from the `oradata` subdirectory of each Oracle mount point set up on the server. Therefore, the process of converting an Oracle installation from `sdisk` to SDD has two parts:

- Change the Oracle mount points' physical devices in `/etc/fstab` from `sdisk` device partition links to the SDD device partition links that access the same physical partitions.
- Re-create any links to raw `sdisk` device links to point to raw SDD device links that access the same physical partitions.

Converting an Oracle installation from `sd` to SDD partitions: 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/rdisk/c1t4d0s4`).

3. Launch the `sqlplus` utility.

4. Type 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. Issue the **ls -l** command on each device link found in step 4 and extract the link source device file name. For example, if you type the command:

```
# ls -l /dev/rdisk/c1t1d0s4
```

A message similar to the following message is displayed:

```
/dev/rdisk/c1t1d0s4 /devices/pci@1f,0/pci@1/scsi@2/sd@1,0:e
```

6. Write down the file ownership and permissions by issuing the **ls -lL** command on either the files in `/dev/` or `/devices` (it yields the same result). For example, if you type the command:

```
# ls -lL /dev/rdisk/c1t1d0s4
```

A message similar to the following message is displayed:

```
crw-r--r-- oracle dba 32,252 Nov 16 11:49 /dev/rdisk/c1t1d0s4
```

7. Complete the basic installation steps in the "Installing SDD" on page 144 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.

Solstice DiskSuite

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.

Installing Solstice DiskSuite for the first time

Perform the following steps if you are installing Solstice DiskSuite on the multiport subsystem server for the first time. Perform the following steps to install the Solstice DiskSuite for the first time on an SDD system:

1. Install SDD using the procedure in the “Installing SDD” on page 144 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 Solstice DiskSuite packages and the answerbook. Do not restart yet.

Note: Do not install the DiskSuite Tool (metatool).

4. Determine which vpath devices you will use to create Disk Suite 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, type the **showvpath** command and press Enter. Reserve at least three partitions of three cylinders each for use as DiskSuite 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 *DiskSuite 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 `/etc/opt/SUNWmd/md.tab` file.

Installing SDD on a system that already has Solaris DiskSuite in place

Perform the following steps if Solaris DiskSuite is already installed:

1. Back up all data.

2. Back up the current Solstice configuration by making a copy of the `/etc/opt/SUNWmd/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 `md.tab`, and that they all come up properly after a restart.
3. Install SDD using the procedure in the “Installing SDD” on page 144 section, if you have not already done so. After the installation completes, type 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 the `/etc/opt/SUNWmd` directory.
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, reverse the procedures in step 6, reinstall the original `md.tab` into `/etc/opt/SUNWmd`, issue the **pkgrm IBMsdd** command, and restart.

8. Stop all applications using DiskSuite, including file systems.
9. Type the following commands for each existing metadevice:


```
metaclear <device>
metainit -a
```
10. Restart your applications.

Setting up 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 144 section, if you have not already done so.
2. Determine which vpath (vpathNs) volumes 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 Solstice DiskSuite if you have not already done so.
5. Create the metatrans device using **metainit**. For example, assume `/dev/dsk/vpath0d` is your UFS master device used in step 3, `/dev/dsk/vpath0e` is its corresponding log device, and `d0` is the trans device that you want to create for UFS logging. Type `metainit d0 -t vpath0d vpath0e` 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 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 /etc/opt/SUNWmd/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 144 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. Type **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 /etc/opt/SUNWmd/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 /etc/opt/SUNWmd/md.tab file and restart the system. Review your steps and try again.

Chapter 7. Using SDD on a Windows NT host system

Attention: SDD does not support Windows NT clustering for systems attached to a SAN Volume Controller.

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 to a SAN Volume Controller. For updated and additional information that is not included in this chapter, see the README on the compact disc or visit the SDD Web site at

www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

and 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

- ESS or SAN Volume Controller or both
- Host system
- For ESS devices: SCSI adapters and cables
- Fibre Channel adapters and cables

Software

- Windows NT 4.0 operating system
- 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 devices

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.storage.ibm.com/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 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 ESS 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.storage.ibm.com/hardsoft/products/ess/supserver.htm

Preparing for SDD installation

Before you install SDD, you must configure the ESS or SAN Volume Controller 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 information about configuring your SAN Volume Controller, refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller 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 LP8000 with the full port driver. When you configure Emulex LP8000 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.

Refer to the *IBM TotalStorage Virtualization Family: SAN Volume Controller Host Systems Attachment Guide* for more information about installing and configuring Fibre Channel adapters for your Windows NT host systems for the SAN Volume Controller.

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 183 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 ESS 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. Type `datapath query adapter` and press Enter. The output includes information about all the installed adapters. In the example shown in the following output, one SCSI 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. Type `datapath query device` and press Enter. In the example shown on page 170, SDD displays 10 ESS devices. There are five physical drives, and one partition has been assigned on each drive for this configuration. Each SDD ESS 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 ESS 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 ESS or SAN Volume Controller.
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 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 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 ESS devices that were installed. In the example shown in the following output, the output includes information about the new SCSI 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 ESS 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.

SAN Volume Controller Preferred Node path-selection algorithm

The SAN Volume Controller is a two-controller disk subsystem. SDD distinguishes the paths to a SAN Volume Controller 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 serial number.

```
Total Devices : 4
DEV#: 0 DEVICE NAME: Disk3 Part0 TYPE: 2145 POLICY: OPTIMIZED
SERIAL: 60056768018000000800000000000001
=====
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: 60056768018000000800000000000001
=====
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: 6005676801800000080000000000009
=====
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: 6005676801800000080000000000009
=====
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 176 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 167 for instructions.)

Adding or modifying a multipath storage configuration to the ESS or SAN Volume Controller

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 ESS 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. Type `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  188    0    10    10
  1  Scsi Port7 Bus0  NORMAL  ACTIVE  204    0    10    10

```

- Type 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:

- Install any additional hardware to the ESS.
- Configure the new storage to the server.
- Restart the Windows NT server. Restarting will ensure correct multipath access to both existing and new storage and to your Windows NT server.
- Verify that the new storage is added correctly. See “Verifying new storage is installed correctly”.

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:

- Click **Start** → **Program** → **Subsystem Device Driver** → **Subsystem Device Driver Management**. An MS-DOS window opens.
- Type 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
```

- Type 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 167 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 10, “The SDD server daemon”, on page 215 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

You can change to a different port number for the SDD server.

Note: Before changing the port number, you must stop the SDD server if it is running. Go to “Stopping the SDD server” on page 190 for instructions. The SDD server will not start on a port number if another application is already using the same port number. Use a port number that is not being used by other applications.

Perform the following steps:

1. Click **Start** → **Settings** → **Control Panel**.
2. Double-click **Services**.
3. Click **Look for SDD_Service**.

4. In the startup parameters: enter `pXXXX`
5. Where `XXXX` is the new port number. For example: `p22000`
6. Click **Start**.

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 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 1.3.4.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 an ESS LUN or to a SAN Volume Controller LUN.

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 path” on page 225 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 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 1.3.3.1 (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 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.

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 167.
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 167.
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.

More information about configuring a Windows NT cluster can be found at :
www.microsoft.com/technet/default.asp

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 be the same for both cluster nodes.

Chapter 8. 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 an ESS device or to a SAN Volume Controller device.

For updated and additional information not included in this chapter, see the README on the compact disc or visit the SDD Web site at

www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

and 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

- ESS or SAN Volume Controller or both
- 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
- 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, make sure that you configure the ESS devices as IBM 2105xxx (where xxx is the ESS model number) on your Windows 2000 host system.

SAN Volume Controller requirements

To successfully install SDD, make sure that you configure the SAN Volume Controller devices as IBM 2145 Fibre Channel-attached devices 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.storage.ibm.com/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 ESS configuration, the number of Fibre Channel adapters attached should be less than or equal to $32 / (n * m)$, where n is the number of ESS ports and m is the number of paths that have access to the ESS from the fabric.
- 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 2000 host system, go to the following Web site at:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Preparing for SDD 1.4.0.0 (or later) installation

Before installing SDD 1.4.0.0 (or later), you must:

Note: If you currently have SDD 1.3.x.x running, then IBM recommends an upgrade to 1.4.0.0 (or later). To upgrade to SDD 1.4.0.0 (or later), see "Upgrading SDD" on page 188.

1. Ensure that all hardware and software requirements are met before you install the SDD. See “Verifying the hardware and software requirements” on page 183 for more information.
2. Configure the ESS to your host system. See “Configuring the ESS or SAN Volume Controller” 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” for more information.

Configuring the ESS or SAN Volume Controller

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*.

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 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.storage.ibm.com/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

SDD supports the Emulex LP8000/LP9002 with full port driver. When you configure the Emulex LP8000/LP9002 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 1.4.0.0 (or later)

The following section describes how to install SDD 1.4.0.0 (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 Microsoft's 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 1.4.0.0 (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 install:
 - a. Shut down your Windows 2000 host system.
 - b. Reconnect all cables that connect the host bus adapters and the ESS 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 5 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 5. 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” on page 188.

Displaying the current version of SDD

You can display the current version of SDD on a Windows 2000 host system by viewing the sdbus.sys file properties. Perform the following steps to view the properties of sdbus.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 sdbus.sys file, and then click **Properties**. The sdbus.sys properties window opens.
4. In the sdbus.sys properties window, click **Version**. The file version and copyright information about the sdbus.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 1.4.0.0 (or later)” on page 186 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”.

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 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 node B.

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 there is 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 1.4.0.0 (or later) immediately after performing a system restart to avoid any potential data loss. Go to “Installing SDD 1.4.0.0 (or later)” on page 186 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 to an ESS 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 through the ESS Storwatch Specialist
 - d. Remove multiple ESS port access if applicable
3. Uninstall SDD. See “Uninstalling SDD” on page 188 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 10, “The SDD server daemon”, on page 215 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 sddsrvc:

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

You can change to a different port number for the SDD server.

Note: You must stop the SDD server if it is running before changing the port number. Go to “Stopping the SDD server” for instructions. The SDD server will not start on a port number if another application is already using the same port number. Use a port number that is not being used by other applications.

Perform the following steps:

1. Click **Start** → **Programs** → **Administrative Tools** → **Computer Management**.
2. Expand the Services and Applications tree.
3. Click **Services**.
4. Right-click **SDD_Service**.
5. Click **Properties**.
6. In startup Parameters: Enter pXXXX, where XXXX is the new port number. For example: p22000
7. Click **Start**.

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. Type `datapath query adapter` and press Enter. The output includes information about all the installed adapters. In the example shown in the following output, one SCSI 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. Type `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. 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 SCSI 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. Type datapath query device and press Enter. The output should include information about any additional devices that were installed. In this example, the output includes information about the new SCSI 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
```

SAN Volume Controller Preferred Node path-selection algorithm

The SAN Volume Controller is a two-controller disk subsystem. SDD distinguishes the paths to a SAN Volume Controller 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 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 an ESS LUN or to a SAN Volume Controller 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 *multiple-path 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 path” on page 225 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 web site for recommended QLogic BIOS settings:

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:

<i>server_1</i>	represents the first server with two host bus adapters (HBAs).
<i>server_2</i>	represents the second server with two HBAs.
<i>hba_a</i>	represents the first HBA for <i>server_1</i>
<i>hba_b</i>	represents the second HBA for <i>server_1</i>
<i>hba_c</i>	represents the first HBA for <i>server_2</i>
<i>hba_d</i>	represents the second HBA for <i>server_2</i>

Perform the following steps to configure a Windows 2000 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 on page 197.

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 1.4.0.0 (or later)” on page 186 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 1.4.0.0 (or later)” on page 186.
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 9. 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 an ESS or to a SAN Volume Controller device.

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

www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

and 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

- ESS or SAN Volume Controller
- 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 2105xxx (where xxx is the ESS model number) on your Windows Server 2003 host system.

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.storage.ibm.com/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.storage.ibm.com/hardsoft/products/ess/supserver.htm

Preparing for SDD 1.4.0.0 (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 1.4.0.0
2. Upgrade Windows 2000 server to Windows Server 2003

Before installing SDD 1.4.0.0 (or later), you must:

1. Ensure that all hardware and software requirements are met before you install the SDD. See “Verifying the hardware and software requirements” on page 199 for more information.
2. Configure the ESS to your host system. See “Configuring the ESS” on page 201 for more information.
3. Configure the Fibre Channel adapters that are attached to your Windows Server 2003 host system. See “Configuring Fibre Channel adapters” on page 201 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 204
 - b. “Installing SDD 1.4.0.0 (or later)”

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*.

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 LP8000/LP9002 with full-port driver. When you configure the Emulex LP8000/LP9002 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 1.4.0.0 (or later)

The following section describes how to install SDD 1.4.0.0 (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 Microsoft's 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 1.4.0.0 (or later) on your system:

1. If this is a fresh install, ensure there is a single connection from your host to your storage. Multipath access to the storage can be added once 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 6, 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 6. 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”.

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 176.
2. Shut down the system.
3. Disconnect all cables that allow the Windows NT host to access to ESS or SAN Volume Controller devices.
4. Restart the system.
5. Perform the Windows NT to Windows Server 2003 upgrade according to your migration plans.
6. Once your host upgrade is complete, install Windows Server 2003-supported HBA drivers.
7. Enable a single path access from your server to the ESS or SAN Volume Controller.
8. Restart your host.
9. Install the latest version of SDD for Windows 2003. See “Installing SDD 1.4.0.0 (or later)” on page 201.
10. Reboot the system, enabling additional paths to the ESS or SAN Volume Controller.

Uninstalling SDD

Attention:

- You must install SDD 1.4.0.0 (or later) immediately after performing a system restart to avoid any potential data loss. Go to “Installing SDD 1.4.0.0 (or later)” on page 201 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 once 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 1.4.0.0 (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 through the ESS Storwatch Specialist
 - d. Remove multiple ESS port access if applicable
3. Uninstall SDD. See “Uninstalling SDD” on page 204 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 1.4.0.0 (or later). This component consists of a Windows application daemon that is installed in addition to the SDD device driver. See Chapter 10, “The SDD server daemon”, on page 215 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

You can change to a different port number for the SDD server.

Note: You must stop the SDD server if it is running before changing the port number. Go to “Stopping the SDD server” for instructions. The SDD server will not start on a port number if another application is already using the same port number. Use a port number that is not being used by other applications.

Perform the following steps:

1. Click **Start** → **Programs** → **Administrative Tools** → **Computer Management**.
2. Expand the Services and Applications tree.
3. Click **Services**.
4. Right-click **SDD_Service**.
5. Click **Properties**.
6. In startup Parameters: Enter pXXXX, where XXXX is the new port number. For example: p22000
7. Click **Start**.

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. Type datapath query adapter and press Enter. The output includes information about all the installed adapters. In the example shown in the following output, one SCSI 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. Type datapath query device 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” on page 208.

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. 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 SCSI 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. Type `datapath query device` and press Enter. The output should include information about any additional devices that were installed. In this example, the output includes information about the new SCSI 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 1.4.0.0 (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 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 1.4.0.0 (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.)

multipath mode

The host system has multiple paths that are configured to an ESS LUN or to a SAN Volume Controller 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 *multiple-path 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 path” on page 225 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:

publibfp.boulder.ibm.com/epubs/pdf/f2bhs00.pdf

Support for Windows Server 2003 clustering

SDD 1.4.0.0 or higher is required to support Windows Server 2003 clustering. SDD 1.4.0.0 or higher 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.

Configuring a Windows Server 2003 cluster on ESS

Information about installing a Windows Server 2003 cluster can be found at : www.microsoft.com/technet/default.asp

Look for product information about Windows Server 2003 clustering.

Configure Windows 2003 cluster with SDD installed

The following variables are used in this procedure:

<i>server_1</i>	represents the first server with two host bus adapters (HBAs).
<i>server_2</i>	represents the second server with two HBAs.
<i>hba_a</i>	represents the first HBA for <i>server_1</i>
<i>hba_b</i>	represents the second HBA for <i>server_1</i>
<i>hba_c</i>	represents the first HBA for <i>server_2</i>
<i>hba_d</i>	represents the second HBA for <i>server_2</i>

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 1.4.0.0 (or later)" on page 186 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 1.4.0.0 (or later)” on page 186.
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. The SDD server daemon

The SDD server (sddsrsv) consists of a UNIX application daemon that is installed in addition to the SDD device driver. The sddsrsv daemon starts automatically after the SDD driver package is installed. The sddsrsv daemon runs in the background at all times. The daemon scans to find failing paths (INVALID, CLOSE_DEAD, or DEAD) at regular intervals between 2-5 minutes. The daemon probes idle paths that are in the CLOSE or OPEN states at regular, 1-minute intervals.

Understanding how the SDD server daemon works

The sddsrsv daemon provides path reclamation and path probing.

Path reclamation

The SDD server will regularly test and recover broken paths that have become operational. It will test INVALID, CLOSE_DEAD, or DEAD paths and detect if these paths have become operational. The daemon sleeps for 2- to 5-minute intervals between consecutive executions. If the test succeeds, then sddsrsv reclaims these paths and changes the states of these paths according to the following characteristics:

- If the state of the vpath is OPEN, then sddsrsv changes the states of INVALID and CLOSE_DEAD paths of that vpath to OPEN.
- If the state of the vpath is CLOSE, then sddsrsv changes the states of CLOSE_DEAD paths of that vpath to CLOSE.
- Sddsrsv changes the states of DEAD paths to OPEN.

Path probing

The SDD server will regularly test CLOSE paths and OPEN paths that are idle to see if they are operational or have become unoperational. The daemon sleeps for 1-minute intervals between consecutive executions. If the test fails, sddsrsv 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 sddsrsv 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 sddsrsv changes the state of the path from CLOSE to CLOSE_DEAD.
- If all paths are not operational, then the sddsrsv sets the last path to the DEAD state.

Note: Sddsrsv will not test those paths which are manually placed offline.

Sddsrsv and StorWatch Expert

Sddsrsv is running on a TCP/IP port. If you are using StorWatch Expert, then it needs to communicate with sddsrsv through a TCP/IP socket on the port on which sddsrsv is running. You should apply your corporate security rules around this port.

Chapter 11. Using the datapath commands

SDD provides commands that you can use to display the status of adapters that are used to access managed devices, or to display the status of devices that the device driver manages. You can also set individual path conditions either to online or offline, or you can set all paths that are connected to an adapter either to online or offline. This chapter includes descriptions of these commands. Table 31 provides an alphabetical list of these commands, a brief description, and where to go in this chapter for more information.

Table 31. Commands

Command	Description	Page
datapath query adapter	Displays information about adapters	218
datapath query adaptstats	Displays performance information for all SCSI and FCS adapters that are attached to SDD devices	219
datapath query device	Displays information about devices	220
datapath query devstats	Displays performance information for a single SDD device or all SDD devices	222
datapath set adapter	Sets all device paths that are attached to an adapter to online or offline	224
datapath set path	Sets the path of a device to online or offline	225
datapath open path	Dynamically opens a path that is in an Invalid state.	226
datapath set device policy	Dynamically changes the path-selection policy of the SDD device or devices	228

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 adapter number for which you want information displayed. If you do not enter an adapter number, information about all adapters is displayed.

Examples

If you type 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.

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 are not functioning.

Failed The adapter is no longer being used by SDD.

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 FCS adapters that are attached to SDD devices. If you do not enter a device number, information about all devices is displayed.

Syntax

▶—datapath query adaptstats—*adapter number*————▶

Parameters

adapter number

The adapter number for which you want information displayed. If you do not enter an adapter number, information about all adapters is displayed.

Examples

If you type 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.

Syntax

```
▶—datapath query device [ device number ] —▶
                        [ -d device model ]
```

Parameters

device number

The device number refers to the device *index* number, 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).

2105E Display all 2105 E models (ESS).

2105F Display all 2105 F models (ESS).

2105F10

Display all 2105 F10 models (ESS).

21058

Display all 2105 800 models (ESS).

2145 Display all 2145 models (SAN Volume Controller).

Examples

If you type the `datapath query device 0` command, the following output displays:

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:

```
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.

Name The name of this device.

Type The device product ID from inquiry data.

Serial The logical unit number (LUN) for this device.

Path The path number.

Adapter

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

Hard Disk

The name of the logical device to which the path is bound.

State The condition of the named device:

Open Path is in use.

Close Path is not being used.

Close_Dead

Path is broken and not being used.

Dead Path is no longer being used. It was either removed by SDD due to errors or manually removed using the **datapath set device M path N offline** or **datapath set adapter N offline** command.

Invalid

The path failed to open.

Mode The mode of the named path. The mode can be either Normal or Offline.

Select The number of times this path was selected for input or output.

Errors The number of input and output errors on a path that is attached to this device.

datapath query devstats

The **datapath query devstats** command displays performance information for a single SDD device or all SDD devices. If you do not enter a device number, information about all devices displays.

Syntax

```
▶▶—datapath query devstats [ device number ] [ -d device model ] ▶▶
```

Parameters

device number

The device number refers to the device *index* number, rather than the SDD device number.

-d *device model*²

The device model 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).

2105E Display all 2105 E models (ESS).

2105F Display all 2105 F models (ESS).

2105F10

Display all 2105 F10 models (ESS).

21058

Display all 2105 800 models (ESS).

2145 Display all 2145 models (SAN Volume Controller).

Examples

If you type the `datapath query devstats 0` command, the following output is displayed:

```
Device #: 0
=====
                Total Read  Total Write  Active Read  Active Write  Maximum
I/O:                387      24502563           0           0         62
SECTOR:              9738      448308668           0           0        2098

Transfer Size:      <= 512      <= 4k      <= 16K      <= 64K      > 64K
                   4355850      1024164      19121140      1665        130

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

The terms used in the output are defined as follows:

Total Read

- I/O: total number of completed Read requests
- SECTOR: total number of sectors that have been read

2. The option to specify a device model cannot be used when you query a specific device number.

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 set adapter

The **datapath set adapter** command sets all device paths attached to an adapter either to online or offline. SDD reserves the last path of a device. The **datapath set adapter offline** command will not remove the last path to a device. 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 ESS or SAN Volume Controller ports through a switch, the **datapath set adapter 0 offline** command fails all the paths that are not set offline.

Attention: The **datapath set adapter offline** command can be issued even when the device is closed.

Syntax

```
▶▶—datapath set adapter—adapter number—online  
offline—▶▶
```

Parameters

adapter number

The adapter number that you want to change.

online

Sets the adapter online.

offline

Sets the adapter offline.

Examples

If you type the `datapath set adapter 0 offline` command:

- Adapter 0 changes to Offline mode and its state changes 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.

datapath set path

The **datapath set 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 path** command can be issued even when the device is closed.

Syntax

```
▶▶—datapath set device—device number path number [ online | offline ] ▶▶
```

Parameters

device number

The device index number that you want to change.

path number

The path number that you want to change.

online

Sets the path online.

offline

Removes the path from service.

Examples

If you type the `datapath set device 0 path 0 offline` command, path 0 for device 0 changes to Offline mode.

datapath open path

The **datapath open 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, rather than the SDD device number.

path number

The path number that you want to change.

Examples

If you type the `datapath query device 8` command, the following output displays:

```
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 type 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:

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 is not being used.

Dead Path is no longer being used.

Invalid

The path failed to open.

Mode The mode of the named path, which is either Normal or Offline.

Select The number of times this path was selected for input and output.

Errors The number of input errors and output errors that are on this path.

datapath set device policy

The **datapath set device policy** command dynamically changes the path selection policy per SDD vpath device.

Syntax

```
▶▶—datapath set device policy —┬─ device number —————▶▶  
                               └─ device number1 device number2 ┘
```

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
- **df**, where *df* indicates the default policy, which is load balancing
- **fo**, where *fo* indicates failover policy

The **rr**, **lb**, and **df** options have the same effect.

Note: You can issue 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.

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.

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
eServer
ES/9000
ESCON
FICON
FlashCopy
HACMP/6000
IBM
IBM logo
IBM TotalStorage
iSeries
MVS/ESA
Netfinity
NetVista
NUMA-Q
Operating System/400
OS/390
OS/400
pSeries
RS/6000
S/390
Seascape
SNAPSHOT

SP
StorWatch
System/360
System/370
System/390
System/400
TotalStorage
Versatile Storage Server
VM/ESA
VSE/ESA
xSeries
zSeries
z/Architecture

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.

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission (FCC) statement

This equipment has been tested and complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the users authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conform à la norme NMB-003 du Canada.

European community compliance statement

This product is in conformity with the protection requirements of EC Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

Germany only

Zulassungsbescheinigung laut Gesetz ueber die elektromagnetische Vertraeglichkeit von Geraeten (EMVG) vom 30. August 1995.

Dieses Geraet ist berechtigt, in Uebereinstimmung mit dem deutschen EMVG das EG-Konformitaetszeichen - CE - zu fuehren.

Der Aussteller der Konformitaetserklaeung ist die IBM Deutschland.

Informationen in Hinsicht EMVG Paragraph 3 Abs. (2) 2:

Das Geraet erfuehlt die Schutzanforderungen nach EN 50082-1 und EN 55022 Klasse A.

EN 55022 Klasse A Geraete beduerfen folgender Hinweise:

Nach dem EMVG:

"Geraete duerfen an Orten, fuer die sie nicht ausreichend entstoert sind, nur mit besonderer Genehmigung des Bundesministeriums fuer Post und Telekommunikation oder des Bundesamtes fuer Post und Telekommunikation betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Stoerungen zu erwarten sind." (Auszug aus dem EMVG, Paragraph 3, Abs.4)

Dieses Genehmigungsverfahren ist nach Paragraph 9 EMVG in Verbindung mit der entsprechenden Kostenverordnung (Amtsblatt 14/93) kostenpflichtig.

Nach der EN 55022:

"Dies ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funkstoerungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Massnahmen durchzufuehren und dafuer aufzukommen."

Anmerkung:

Um die Einhaltung des EMVG sicherzustellen, sind die Geraete wie in den Handbuechern angegeben zu installieren und zu betreiben.

Japanese Voluntary Control Council for Interference (VCCI) class A statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean government Ministry of Communication (MOC) statement

Please note that this device has been approved for business purposes with regard to electromagnetic interference. If you find this is not suitable for your use, you may exchange it for a non-business purpose one.

Chinese Class A warning statement

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

中华人民共和国“A类”警告声明

声 明

此为A类产品，在生活环境中，该产品可能会造成无线电干扰。在这种情况下，可能需要用户对其干扰采取切实可行的措施。

Taiwan class A compliance statement

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

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 Agreement³ 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:

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.

3. Form Z125-4144

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 SC1.

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*.

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 Seascapes 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). Formerly called IBM StorWatch Enterprise Storage Server 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-stripping methodology.

nonremovable medium. A recording medium that cannot be added to or removed from a storage device.

nonvolatile storage (NVS). In the ESS, memory that stores active write data to avoid data loss in the event of a power loss.

NVS. See *nonvolatile storage*.

O

octet. In Internet Protocol (IP) addressing, one of the four parts of a 32-bit integer presented in dotted decimal notation. dotted decimal notation consists of four 8-bit numbers written in base 10. For example, 9.113.76.250 is an IP address containing the octets 9, 113, 76, and 250.

OEMI. See *original equipment manufacturer's information*.

open system. A system whose characteristics comply with standards made available throughout the industry and that therefore can be connected to other systems complying with the same standards. Applied to the ESS, such systems are those hosts that connect to the ESS through SCSI or FCP protocols. See also *small computer system interface* and *fibre-channel protocol*.

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 @serverfamilies 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 AS/400, a disk storage device that is protected from data loss by RAID techniques. An AS/400 host does not mirror a volume configured as a protected volume, while it does mirror all volumes configured as unprotected volumes. The 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 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*.

Seascope 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 2⁴⁰ or 1 099 511 627 776 bytes.

terminal emulator. In the ESS, a function of the ESS Master Console that allows it to emulate a terminal.

thousands of power-on hours (KPOH). A unit of time used to measure the mean time between failures (MTBF).

time sharing option (TSO). An operating system option that provides interactive time sharing from remote terminals.

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 164

A

about this book xv

accessing

AIX

add a data path volume to a volume group SMIT panel 74

add a volume group with data path devices SMIT panel 73

add paths to available data path devices SMIT panel 72

back up a volume group with data path devices SMIT panel 75

configure a defined data path device SMIT panel 73

define and configure all data path devices SMIT panel 72

display data path device adapter status SMIT panel 72

display data path device configuration SMIT panel 70

Display Data Path Device Status SMIT panel 71

remake a volume group with data path devices SMIT panel 75

remove a data path device SMIT panel 73

Remove a Physical Volume from a Volume Group SMIT panel 74

adapter

firmware level 13

upgrading firmware level to (sf320A9) 14

adapters 134

configuring

Linux 118

Novell 135

Windows 2000 185

Windows NT 167

Windows Server 2003 201

LP70000E 9

adding

paths

Windows 2000 host systems 190

Windows NT 168

Windows Server 2003 host systems 206

storage for Windows NT host systems 174, 175

adding paths

AIX

to SDD devices volume group 35

from AIX 4.3.2 volume group 35

addpaths

utility programs, AIX 76

addpaths command 35, 67, 76

agreement for licensed internal code 233

AIX

4.3.2

32-bit applications 15

64-bit applications 15

4.3.3

32-bit applications 15

64-bit applications 15

5.1.0

32-bit applications 15

64-bit applications 15

accessing

add a data path volume to a volume group SMIT panel 74

add a volume group with data path devices SMIT panel 73

add paths to available data path devices SMIT panel 72

back up a volume group with data path devices SMIT panel 75

configure a defined data path device SMIT panel 73

define and configure all data path devices SMIT panel 72

display data path device adapter status SMIT panel 72

display data path device configuration SMIT panel 70

Display Data Path Device Status SMIT panel 71

remake a volume group with data path devices SMIT panel 75

Remove a Physical Volume from a Volume Group SMIT panel 74

removing a data path device SMIT panel 73

adding paths 35

applications

32-bit 15

64-bit 15

backing-up files belonging to an SDD volume group 69

changing the path-selection policy 34, 126

configuring

volume group for failover protection 61

configuring SDD 24

dynamically adding paths to SDD devices of a volume group 35

error log messages 87

new and modified messages by SDD for HACMP environment only 88

exporting

volume group with SDD 63

extending

an existing SDD volume group 68

importing

volume group with SDD 62

installing SDD 18

migrating

an existing non-SDD volume group to SDD vpath devices in concurrent mode 85

AIX (*continued*)

- migrating (*continued*)
 - non-SDD volume group to a SAN Volume Controller SDD multipath volume group in concurrent mode 83
 - non-SDD volume group to an ESS SDD multipath volume group in concurrent mode 83
- recovering
 - from mixed volume groups 68
- removing SDD from a host system 33
- restoring files belonging to an SDD volume group 69
- SDD utility programs 76
- SDD-specific SMIT panels 69
- the loss of a device path 65
- unconfiguring SDD 25
- unsupported environments 8
- verifying SDD 30
- verifying SDD installation 20

AIX 5.1.0

- 32-bit 15
- 64-bit 15

AIX 5.2.0

- 32-bit 15
- 64-bit 15

AIX Fibre Channel requirements 11

AIX host system

- disk driver 1
- protocol stack 1

AIX levels

- AIX 4.3.2
 - required PTFs 8
- AIX 4.3.3
 - required maintenance level 8

AIX trace 87

article

- Microsoft Knowledge Base Article Number Q293778
 - information about removing multipath access to your shared volume 177

B

- backing-up AIX files belonging to an SDD volume group 69
- BIOS, disabling 167, 185, 201
- block disk device interfaces (SDD) 90, 143
- boot -r command 162
- bootinfo -K command 15

C

- Canadian compliance statement 231
- cat /proc/modules command 121
- cat /proc/scsi/scsi command 122
- cat /proc/scsi/xxx/N command 122
- cat /proc/sdd command 122
- cd /mnt command 119
- cd /opt/IBMsdd command 120
- cd /opt/IBMsdd/bin command 120
- cfallvpath 17

- cfgmgr
 - run *n* times where *n* represents the number of paths per SDD device. 67
 - run for each installed SCSI or fibre adapter 67
- cfgmgr command 13, 36, 67
- cfgvpath command 124
- changing
 - path-selection policy for AIX 34, 126
 - path-selection policy for HP 99
 - path-selection policy for SUN 152
 - SDD hardware configuration
 - HP host systems 98
 - Sun hosts 150
 - to the /dev directory
 - HP host systems 108
- chdev command 66
- chgrp command 116, 161
- chkconfig - -level X sdd on command 125
- chkconfig - -level X sdd off command 125
- chkconfig - -list sdd command 125
- chkvpenv command 121
- chmod command 116, 161
- chown command 116
- class A compliance statement, Taiwan 233
- command
 - addpaths 35, 67, 76
 - bootinfo -K 15
 - cfgmgr 13, 67
 - running *n* times for *n*-path configurations 36, 67
 - running for each relevant SCSI or FCP adapter 36
 - chdev 66
 - datapath open device 226
 - datapath query adapter 218
 - datapath query adaptstats 219
 - datapath query device 61, 67, 220
 - datapath query devstats 222
 - datapath set adapter 224
 - datapath set device 225
 - datapath set device 0 path 0 offline 225
 - datapath set device N policy rr/fo/lb/df 35, 99, 126, 152
 - datapath set device policy 228
 - dpovgfix 61, 77
 - dpovgfix vg-name 36, 66
 - extendvg 69
 - extendvg4vp 69, 78
 - hd2vp and vp2hd 76
 - HP host system
 - hd2vp 98
 - vgexport 105, 106
 - vgimport 106
 - vp2hd 98
 - installp 12
 - instfix -i | grep IY10201 12
 - instfix -i | grep IY10994 12
 - instfix -i | grep IY11245 12
 - instfix -i | grep IY13736 12
 - instfix -i | grep IY17902 12
 - instfix -i | grep IY18070 12
 - ls -al /unix 15

command (*continued*)

- lscfg -vl fcsN 13
- lsdev -Cc disk 13
- lsdev -Cc disk | grep 2105 24
- lsdev -Cc disk | grep SAN Volume Controller 24
- lslpp -l ibmSdd_432.rte 20
- lslpp -l ibmSdd_433.rte 20
- lslpp -l ibmSdd_510.rte 20
- lslpp -l ibmSdd_510nchacmp.rte 20, 21
- lspv 36, 64
- lsvg -p vg-name 62
- lsvpcfg 36, 60, 65, 77
- mkdev -l vpathN 25
- mksysb restore command 65
- mkvg 62
- mkvg4vp 62, 78
- odmget -q "name = ioaccess" CuAt 49
- querysn 16, 78
- restvg 69
- restvg4vp 69
- rmdev 67, 68
- rmdev -dl dpo -R 33
- rmdev -dl fcsN -R 13
- rmdev -l dpo -R 25
- savevg 69
- savevg4vp 69
- shutdown -rF 13
- smitty deinstall 12
- table of, in installation package 17
- using 217
- varyoffvg 24

commands

- /opt/IBMsdd/bin/showvpath 164
- boot -r 162
- cat /proc/modules 121
- cat /proc/scsi/scsi 122
- cat /proc/scsi/xxx/N 122
- cat /proc/sdd 122
- cd /mnt 119
- cd /opt/IBMsdd 120
- cd /opt/IBMsdd/bin 120
- cfgvpath 124
- chgrp 116, 161
- chkconfig - -level X sdd off 125
- chkconfig - -level X sdd on 125
- chkconfig - -list sdd 125
- chkvpenv 121
- chmod 116, 161
- chown 116
- datapath query
 - adapter 181, 197, 213
 - device 124, 168, 187, 188, 197, 202, 213
- datapath query device 94
- datapath set adapter # offline 196, 211
- datapath set adapter offline 196, 211
- esvpcfg 131
- insmod ./vpath.o 121
- ls -l 120
- lsvpcfg 124
- lsvpsd 123
- metadb -a <device> 163

commands (*continued*)

- metadb -d -f <device> 163
- metadb -i 162
- metainit 162
- metainit d <metadevice number> -t <"vpathNs" - master device> <"vpathNs" - logging device> 164
- metastat 162, 164
- newfs 163
- oraInst /m 114, 159
- pkgmgr IBMsdd 163
- rmvpath xxx 124
- rpm -e IBMsdd command 127
- rpm -qi IBMsdd 120, 127
- rpm -ql IBMsdd 120, 127
- sdd restart 125
- sdd start 125
- sdd stop 125
- showvpath 109, 160, 161, 162, 163
- shutdown -i6 -y -g0 163
- umount 164
- umount /cdrom 150
- unmod ./sdd-mod.o 127
- commands datapath set adapter # offline 180
- commands datapath set adapter offline 180
- comments, how to send xxvii
- communications statement 231
- compliance statement
 - German 232
 - radio frequency energy 231
 - Taiwan class A 233
- concurrent download of licensed internal code
 - AIX 37
 - SDD 6
- configuring
 - additional paths on a Windows NT host system 170
 - AIX
 - cabling storage side switch ports 13
 - ESS 10
 - Fibre Channel attached devices 11, 13
 - SAN Volume Controller 11
 - volume group for failover protection 61
 - clusters with SDD
 - Windows 2000 host system 196
 - Windows NT host system 180
 - Windows Server 2003 host system 211
 - ESS
 - HP host systems 90
 - Linux host systems 118
 - Novell host systems 135
 - Sun host systems 142
 - Windows 2000 185
 - Windows NT 167
 - Windows Server 2003 201
 - Fibre Channel adapters
 - Linux host systems 118
 - Novell host systems 135
 - Windows 2000 host system 185
 - Windows NT host system 167
 - Windows Server 2003 host system 201
 - SAN Volume Controller
 - Linux host systems 118

- configuring *(continued)*
 - SAN Volume Controller *(continued)*
 - Sun host systems 142
 - SCSI adapters
 - Windows 2000 host systems 185
 - Windows NT 167
 - Windows Server 2003 host systems 201
 - SDD
 - at system startup 125
 - Linux host systems 120, 121
 - Novell host systems 136
 - SUN host systems 150
 - Windows NT host system 168
 - SDD for AIX host 24
 - SDD on AIX 24
- configuring a vpath device to the Available state 68
- configuring all vpath devices to the Available state 68
- conversion script
 - vp2hd 25
- conversion scripts
 - hd2vp 76
 - vp2hd 76
- creating
 - device node for the logical volume device in an HP host systems 108
 - directory in /dev for the volume group in an HP host systems 108
 - file system on the volume group in an HP host systems 109
 - logical volume in an HP host systems 109
 - new logical volumes in an HP host systems 108
 - physical volume in an HP host systems 109
 - volume group in an HP host systems 109
- customizing
 - Network File System file server 112
 - Oracle 113, 158
 - standard UNIX applications 107, 157

D

- database managers (DBMS) 143
- datapath
 - open device command 226
 - query
 - adapter command 181, 197, 213
 - device command 124, 168, 181, 197, 213
 - query adapter command 218
 - query adaptstats command 219
 - query device command 220
 - query devstats command 222
 - query set adapter command 224
 - set adapter # offline command 180, 196, 211
 - set adapter offline command 180, 196, 211
 - set device command 225
 - set device policy command 228
- datapath query adapter 100
- datapath query device command 61, 67
- datapath set device 0 path 0 offline command 225
- datapath set device N policy rr/fo/lb/df command 35, 99, 126, 152
- definitions 235

- determining
 - AIX
 - adapter firmware level 13
 - major number of the logical volume device for an HP host systems 108
 - size of the logical volume for an HP host systems 110
 - device driver 142
 - devices.fcp.disk.ibm2105.rte 7
 - devices.fcp.disk.ibm2145.rte 7, 11
 - devices.scsi.disk.ibm2105.rte 7
 - displaying
 - AIX
 - ESS vpath device configuration 60
 - current version of SDD
 - Windows 2000 188
 - Windows Server 2003 203
 - current version of the SDD
 - Windows NT 178
 - dpovgfix command 61, 77
 - dpovgfix vg-name command 36, 66
 - dynamic I/O load-balancing 5

E

- electronic emission notices 231
- enhanced data availability 4
- error log messages
 - AIX
 - new and modified messages by SDD for HACMP environment only 88
 - VPATH_DEVICE_OFFLINE 87
 - VPATH_DEVICE_ONLINE 87
 - VPATH_PATH_OPEN 87
 - VPATH_XBUF_NOMEM 87
- error log messages for ibmSdd_433.rte installation package for SDD
 - AIX
 - VPATH_FAIL_RELPRESERVE 88
 - VPATH_OUT_SERVICE 87
 - VPATH_RESV_CFLICT 88
- error recovery
 - HP-UX 11.0 host system 103
 - HP-UX 11i host system 103
 - SUN host system 153
- error recovery and retry policy
 - algorithms
 - Windows 2000 host systems 194
 - Windows NT host systems 178
 - Windows Server 2003 host systems 209
 - error recovery and retry policy, for Windows 2000 host systems 194
 - error recovery and retry policy, for Windows NT host systems 178
 - error recovery and retry policy, for Windows Server 2003 host systems 209
- error recovery policy
 - algorithms
 - Linux host systems 129
 - error recovery policy, for Linux host systems 129

ESS
AIX
 displaying vpath device configuration 60
 configuring for HP 90
 configuring for Linux 118
 configuring for Novell 135
 configuring for Sun 142
 configuring on
 Windows 2000 185
 Windows Server 2003 201
 configuring on Windows NT 167
ESS devices (hdisks) 81
ESS LUNs 81
European Community Compliance statement 232
exporting a volume group with SDD, AIX 63
extending an existing SDD volume group, AIX 68
extendvg command 69
extendvg4vp command 69, 78

F

failover 5
failover protection, AIX
 creating a volume group from a single-path vpath device 65
 losing 64
 manually deleted devices and running the configuration manager 67
 providing load-balancing and failover protection 60
 side effect of running the disk change method 65
 the loss of a device path 65
 when it doesn't exist 60
Federal Communications Commission (FCC) statement 231
Fibre Channel adapters
 configuring
 Linux host systems 118
 Novell host systems 135
 configuring for Windows 2000 185
 configuring for Windows Server 2003 201
 supported
 HP host systems 89
 Linux host systems 117
 Sun host systems 141
 Windows 2000 host systems 184
 Windows NT host systems 166
 Windows Server 2003 host systems 200
 supported on AIX host systems 9
 supported on Novell host systems 134
Fibre Channel device drivers
 configuring for AIX 11
 devices.common.IBM.fc 11
 devices.fcp.disk 11
 devices.pci.df1000f7 11
 installing for AIX 11
 supported on AIX host systems 9
 supported on Novell host systems 134

G

German compliance statement 232
glossary 235

H

HACMP
 concurrent mode 42
 hd2vp conversion script 49
 node failover 58
 nonconcurrent mode 42
 persistent reserve 49
 recovering paths 58
 SDD persistent reserve attributes 48
 sddsrvc 48
 software support for concurrent mode 46
 software support for nonconcurrent mode 43
 special requirements 49
 supported features 47
hardware configuration
 changing
 HP host systems 98
 Sun host systems 150
hardware requirements
 HP
 host systems 89
 Linux host systems 117
 Sun host systems 141
hd2vp and vp2hd command 76
hd2vp command
 HP host system 98
hdisk device
 chdev 65
 modify attributes 65
High Availability Cluster Multi-Processing (HACMP) 42
HP
 changing the path-selection policy 99
 SCSI disk driver (sdisk) 90
 unsupported environments 89
HP host system
 disk driver 1
 protocol stack 1
HP host systems
 changing
 SDD hardware configuration 98
 to the /dev directory 108
 creating
 a file system on the volume group 109
 a logical volume 109
 a volume group 109
 device node for the logical volume device 108
 directory in /dev for the volume group 108
 new logical volumes 108
 physical volume 109
 determining
 major number of the logical volume 108
 size of the logical volume 110
 installing Oracle 113
 installing SDD 92
 converting an Oracle installation from sdisk 115

- HP host systems *(continued)*
 - installing SDD *(continued)*
 - on a Network File System file server 112
 - on a system that already has Network File System file server 112
 - on a system that already has Oracle 114
 - using a file system 114
 - using raw partitions 115
 - mounting the logical volume 109
 - re-creating
 - existing logical volume 110
 - logical volume 111
 - physical volume 108, 110
 - volume group 111
 - removing
 - existing logical volume 110
 - existing volume group 110
 - logical volumes 110
 - SDD 89
 - setting the correct timeout value for the logical volume manager 111
 - setting up Network File System for the first time 112
 - setting up Oracle
 - using a file system 113
 - using raw partitions 113
 - standard UNIX applications 107
 - understanding how SDD works 90
 - upgrading SDD 90, 96
 - using applications with SDD 107
- HP-UX
 - disk device drivers 94, 107
 - LJFS file system 112
 - operating system 89
- HP-UX 11.0
 - 32-bit 90, 91
 - 64-bit 90, 91
- HP-UX 11i
 - 32-bit 90, 91
 - 64-bit 90, 91

I

- ibm2105.rte 11
- ibm2105.rte ESS package 9, 134
- ibmSdd_433.rte installation package
 - for SDD 1.2.2.0
 - removing 49
 - for SDD 1.3.2.0. vpath devices
 - unconfiguring 49
- importing a volume group with SDD, AIX 62
- Industry Canada Compliance statement 231
- insmod ./sdd-mod.o command 121
- installation package
 - AIX 18
 - devices.sdd.nn.rte 18
 - devices.sdd.43.rte 15, 25
 - devices.sdd.51.rte 15, 25
 - devices.sdd.52.rte 15
 - ibmSdd_432.rte 20, 33, 42, 43, 83
 - ibmSdd_433.rte 20, 33, 43, 49, 63, 83, 87, 88
 - ibmSdd_510.rte 21, 33, 42, 43
 - installation package *(continued)*
 - AIX *(continued)*
 - ibmSdd_510nchacmp.rte 33, 42, 43
 - installing
 - additional paths on a Windows NT host system 170
 - AIX
 - Fibre Channel device drivers 11
 - planning 7
 - SDD 18
 - an Oracle installation from sdisk on an HP host systems 115
 - converting an Oracle installation from sdisk on a Sun host system 161
 - Novell
 - planning 133
 - Oracle
 - HP host systems 113
 - Sun host systems 158
 - SDD
 - HP host systems 89, 92
 - Linux host systems 117, 119
 - Novell host systems 136
 - Sun host systems 141, 144
 - Windows 2000 host system 183
 - Windows NT host system 165, 167
 - Windows Server 2003 host system 199
 - SDD 1.4.0.0
 - Windows Server 2003 host system 201
 - SDD 1.4.0.0 (or later)
 - Windows 2000 host system 186
 - SDD on a Network File System file server on a Sun host system 157
 - SDD on a Network File System file server on an HP host systems 112
 - SDD on a system that already has Network File System file server 158
 - SDD on a system that already has Network File System file server on an HP host systems 112
 - SDD on a system that already has Oracle on a Sun host system 160
 - SDD on a system that already has Oracle on an HP host systems 114
 - SDD on a system that already has Solaris DiskSuite in place on a Sun host system 162
 - Solstice DiskSuite for the first time on a Sun host system 162
 - using a file system on a Sun host system 160
 - using a file system on an HP host systems 114
 - using raw partitions on a Sun host system 160
 - using raw partitions on an HP host systems 115
 - vpath on a system that already has UFS logging in place on a Sun host system 164
 - installp command 12
 - instfix -i | grep IY10201 command 12
 - instfix -i | grep IY10994 command 12
 - instfix -i | grep IY11245 command 12
 - instfix -i | grep IY13736 command 12
 - instfix -i | grep IY17902 command 12
 - instfix -i | grep IY18070 command 12

J

Japanese Voluntary Control Council for Interference
(VCCI) statement 232

K

KB 223
Korean government Ministry of Communication (MOC)
statement 233

L

licensed internal code
 agreement 233
Linux
 unsupported environments 117
Linux host system
 disk driver 1
 protocol stack 1
Linux host systems
 configuring
 ESS 118
 Fibre Channel adapters 118
 SAN Volume Controller 118
 SDD 120, 125
 error recovery policy 129
 installing SDD 117, 119
 maintaining vpath device configuration
 persistence 125
 partitioning SDD vpath devices 130
 preparing SDD installation 118
 removing SDD 127
 upgrading SDD 119
 using
 multiple command for SDD configuration 121
 single command for SDD configuration 125
 standard UNIX applications 131
 verifying SDD installation 120
load-balancing, AIX 60
loading SDD on Linux 120, 121
loading SDD on Novell 136
loading SDD on SUN 150
logical volume manager 143
losing failover protection, AIX 64
ls -al /unix command 15
ls -l command 120
lscfg -vl fcsN command 13
lsdev -Cc disk | grep 2105 command 24
lsdev -Cc disk | grep SAN Volume Controller
 command 24
lsdev -Cc disk command 13
lspp -l **Sdd** command 20
lspp -l *ibmSdd_432.rte* command 20
lspp -l *ibmSdd_433.rte* command 20
lspp -l *ibmSdd_510.rte* command 20
lspp -l *ibmSdd_510nchacmp.rte* command 20, 21
lspv command 36, 64
lsvg -p vg-name command 62
lsvpcfg 17
lsvpcfg command 36, 60, 65, 77, 124, 131

lsvpcfg utility programs, AIX 77
lsvpd command 123

M

maintaining vpath device configuration persistency, for
 Linux host systems 125
Manual Exclusion of ESS devices from the SDD
 configuration 16
 metadb -a <device> command 163
 metadb -d -f <device> command 163
 metadb -i command 162
 metainit command 162
 metainit d <metadevice number> -t <"vpathNs" - master
 device> <"vpathNs" - logging device> command 164
 metastat command 162, 164
migrating
 AIX
 an existing non-SDD volume group to SDD vpath
 devices in concurrent mode 85
 non-SDD volume group to a SAN Volume
 Controller SDD multipath volume group in
 concurrent mode 83
 non-SDD volume group to an ESS SDD multipath
 volume group in concurrent mode 83
mirroring logical volumes 84
mkdev -l vpathN command 25
mksysb restore command 65
mkvg command 62
mkvg4vp command 62, 78
modifying multipath storage configuration to the ESS,
 Windows NT host system 174
mounting the logical volume, HP 109

N

newfs command 163
notices
 electronic emission 231
 European community 232
 FCC statement 231
 German 232
 Industry Canada 231
 Japanese 232
 Korean 233
 licensed internal code 233
 notices statement 229
 Taiwan 233
Novell host systems
 configuring
 ESS 135
 Fibre Channel adapters 135
 SDD 136
 error logging 139
 error reporting 139
 installing SDD 136
 preparing SDD installation 135
 removing SDD 140
 supported environments 133
 unsupported environments 134

O

odmget -q "name = ioaccess" CuAt command 49
orairst /m command 114, 159

P

partitioning SDD vpath devices, for Linux host systems 130
path-failover protection system 5
path-selection policy
 changing 35, 99, 126, 152
 default (optimized) 138
 failover only 34, 99, 126, 138, 152
 load balancing 99, 126, 138, 152
 load-balancing 34
 round robin 34, 99, 126, 137, 152
Persistent Reserve command set 48
pkgm IBMssdd command 163
planning
 AIX
 adapter firmware level 13, 14
 ESS 10
 Fibre Channel attached devices 11, 13
 Fibre Channel device drivers 11
 preparing 10
 SAN Volume Controller 11
 AIX installation 7
 ESS
 HP host systems 90
 Linux host systems 118
 Novell host systems 135
 Sun host systems 142
 Windows 2000 host system 185
 Windows NT host system 167
 Windows Server 2003 host system 201
 Fibre Channel adapters
 Windows 2000 host system 185
 Windows NT host system 167
 Windows Server 2003 host system 201
 hardware and software requirements on a Sun host system 141
 hardware and software requirements on an HP host systems 89
 hardware requirements, AIX
 ESS 7
 Fibre adapters and cables 7
 Host system 7
 SAN Volume Controller 7
 SCSI adapters and cables 7
 hardware requirements, Windows 2000
 ESS 183
 hardware requirements, Windows NT
 ESS 165
 hardware requirements, Windows Server 2003
 ESS 199
 host system requirements, AIX 8
 ESS 9
 Fibre 9
 SAN Volume Controller 9
 SCSI 9

planning (continued)

 host system requirements, Novell 133
 ESS 134
 Fibre 134
 SCSI 134
 host system requirements, Windows 2000
 ESS 183
 host system requirements, Windows NT
 ESS 165
 host system requirements, Windows Server 2003
 ESS 199
 installation of SDD
 HP host systems 90
 Sun host systems 143
 Novell installation 133
 preparing
 Sun host systems 142
 preparing for SDD installation on an HP host systems 90
 SAN Volume Controller
 Linux host systems 118
 Sun host systems 142
 SCSI adapters
 Windows NT host systems 167
 SDD
 HP host systems 89
 Linux host systems 117, 118
 Novell host systems 135
 Sun host systems 141
 Windows 2000 host system 184
 Windows NT host system 165
 Windows Server 2003 host system 200
 software requirements
 Windows 2000 operating system 183
 Windows NT operating system 165
 Windows Server 2003 operating system 199
 software requirements, AIX
 AIX operating system 7
 ibm2105.rte ESS package 7
 SCSI and Fibre Channel device drivers 7
 Windows 2000
 ESS 185
 Windows Server 2003
 ESS 201
 postinstallation of SDD
 HP host systems 94
 Sun host systems 148
 preparing
 AIX
 SDD installation 10
 configure on AIX 24
 SDD
 HP host systems 90
 Linux host systems 118
 Novell host systems 135
 Windows 2000 installation 184
 Windows NT host system 166
 Windows Server 2003 installation 200
 SDD installation
 Sun host systems 142

- providing
 - AIX
 - failover protection 60
 - load-balancing 60
- pSeries 690 with static LPARs configured 42
- pvid 83
- PVID 64

R

- radio frequency energy compliance statement 231
- raw
 - device interface (sd) 143
 - device interface (sdisk) 90
- re-creating
 - existing logical volume
 - on a HP host systems 110
 - physical volume
 - on an HP host systems 108
 - the logical volume
 - on an HP host systems 111
 - the physical volume
 - on a HP host systems 110
 - the volume group
 - on an HP host systems 111
- recovering from mixed volume groups 68
- recovering from mixed volume groups, AIX 68
- recovery procedures for HP 109, 112
- removing
 - existing logical volume
 - on an HP host system 110
 - existing volume group
 - on an HP host systems 110
 - logical volumes
 - on an HP host systems 110
 - SDD
 - Linux host systems 127
 - Novell host systems 140
 - Windows 2000 host system 188
 - Windows NT host system 176
 - Windows Server 2003 host system 204
 - SDD from an AIX host 33
 - SDD from an AIX host system 33
- removing SDD in a two-node cluster environment 189, 205
- Replacing Manually Excluded ESS devices in the SDD configuration 17
- requirements
 - ESS
 - Windows 2000 host system 183
 - Windows NT 165
 - Windows Server 2003 host system 199
 - hardware and software
 - Linux host systems 117
 - hardware and software on a Sun host system 141
 - hardware and software, HP 89
 - hardware, AIX
 - ESS 7
 - Fibre adapters and cables 7
 - Host system 7
 - SAN Volume Controller 7

- requirements (*continued*)
 - hardware, AIX (*continued*)
 - SCSI adapters and cables 7
 - hardware, Windows 2000
 - ESS 183
 - hardware, Windows NT
 - ESS 165
 - hardware, Windows Server 2003
 - ESS 199
 - host system, AIX 8
 - ESS 9
 - Fibre 9
 - SAN Volume Controller 9
 - SCSI 9
 - host system, Novell 133
 - ESS 134
 - Fibre 134
 - SCSI 134
 - software
 - Windows 2000 operating system 183
 - Windows NT operating system 165
 - Windows Server 2003 operating system 199
 - software, AIX
 - AIX operating system 7
 - ibm2105.rte ESS package 7
 - SCSI and Fibre Channel device drivers 7
- restoring
 - AIX
 - files belonging to an SDD volume group 69
- restvg command 69
- restvg4vp command 69
- reviewing the existing SDD configuration information,
 - Windows NT 168, 174
- rmdev -dl dpo -R command 33
- rmdev -dl fcsN -R command 13
- rmdev -l dpo -R command 25
- rmdev command 67, 68
- rmvpath xxx command 124
- rpm -e IBMsd command 127
- rpm -qi IBMsd command 120, 127
- rpm -ql IBMsd command 120, 127

S

- SAN Volume Controller
 - configuring for Linux 118
 - configuring for Sun 142
- SAN Volume Controller Preferred Node path selection
 - algorithm 34, 100, 153, 173, 193
- savevg command 69
- savevg4vp command 69
- SCSI adapter support
 - AIX host systems 9
 - HP host systems 89
 - Novell host systems 134
 - Sun host systems 141
 - Windows 2000 host system 184
 - Windows NT host system 166
 - Windows Server 2003 host system 200
- SCSI-3 Persistent Reserve command set 48

- SDD
 - architecture 1
 - displaying the current version on Windows 2000 188
 - displaying the current version on Windows Server 2003 203
 - how it works on an HP host system 90
 - how it works on Sun 142
 - installation scenarios 90
 - installing
 - AIX 7
 - HP host system 89
 - Linux 117
 - Novell 133
 - Sun host systems 141
 - Windows 2000 host system 183, 186
 - Windows NT 165
 - Windows Server 2003 host system 199, 201
 - introduction 1
 - overview 1
 - postinstallation of SDD
 - HP host systems 94
 - postinstallation on Sun host systems 148
 - removing SDD on Windows NT 176
 - uninstalling
 - HP host systems 97
 - upgrading
 - HP host systems 90
 - Windows 2000 188
 - Windows Server 2003 204
 - using applications
 - with SDD on HP Network File System file server 112
 - with SDD on HP standard UNIX applications 107
 - with SDD on HP, Oracle 113
 - with SDD on Linux standard UNIX applications 131
 - with SDD on Sun Network File System file Server 157
 - with SDD on Sun standard UNIX applications 157
 - with SDD on Sun, Oracle 158
 - verifying additional paths to SDD devices 171, 192, 208
 - verifying configuration 30
 - Web site xvi
- SDD configuration
 - checking 30
- SDD devices
 - reconfiguring 36
- sdd restart command 125
- SDD server daemon 215
 - AIX host system 18, 37, 128
 - HP host system 102
 - SUN host system 155
 - Windows 2000 host system 189
 - Windows NT host system 177
 - Windows Server 2003 host system 205
- sdd start command 125
- sdd stop command 125
- SDD utility programs, AIX 76
- SDD vpath devices 81
- sddserver.rte
 - AIX host system 14
- sddsrv 48, 215
 - AIX host system 18, 37, 128
 - HP host system 102
 - SUN host system 155
 - Windows 2000 host system 189
 - Windows NT host system 177
 - Windows Server 2003 host system 205
- sddsrv for ESS Expert
 - AIX host system 14
- setting up
 - correct timeout value for the logical volume manager on an HP host system 111
 - Network File System for the first time on a Sun host system 157
 - Network File System for the first time on an HP host system 112
 - Oracle using a file system
 - HP host system 113
 - Sun host system 158
 - Oracle using raw partitions
 - HP host system 113
 - Sun host system 158
 - UFS logging on a new system on a Sun host system 163
- showvpath command 109, 160, 161, 162, 163
- shutdown -i6 -y -g0 command 163
- shutdown -rF command 13
- smit
 - definition 19
- SMIT
 - configuring
 - SDD for Windows NT host system 168
- smitty
 - definition 19
- smitty deinstall command 12
- software requirements
 - for SDD on HP 89
 - for SDD on Linux 117
 - for SDD on Sun 141
- Solaris
 - host system
 - upgrading the Subsystem Device Driver on 143
 - operating system
 - upgrading SDD 141
 - sd devices 157
 - UFS file system 157
- statement
 - of compliance
 - Canada 231
 - European 232
 - Federal Communications Commission 231
 - Japan 232
 - Korean government Ministry of Communication (MOC) 233
 - Taiwan 233
- Subsystem device driver, see SDD. 162
- Sun
 - supported environments 141

- Sun (*continued*)
 - unsupported environments 142
- SUN
 - changing the path-selection policy 152
- Sun disk device drivers 143
- Sun host system
 - disk driver 1
 - protocol stack 1
- Sun host systems
 - changing SDD hardware configuration 150
 - installing
 - Solstice DiskSuite for the first time 162
 - vpath on a system that already has UFS logging in place 164
 - installing Oracle 158
 - installing SDD 144
 - converting an Oracle installation from sdisk Network File System file server 161
 - Network File System file server 157
 - system that already has Network File System file server 158
 - system that already has Oracle 160
 - system that already has Solaris DiskSuite in place 162
 - using a file system 160
 - using raw partitions 160
 - Oracle 158
 - SDD 141
 - SDD postinstallation 148
 - setting up
 - Network File System for the first time 157
 - UFS logging on a new system 163
 - Solstice DiskSuite 162
 - standard UNIX applications 157
 - understanding how SDD works 142
 - upgrading SDD 150
 - using applications with SDD 157
- SUN host systems
 - configuring
 - SDD 150
- Sun SCSI disk driver 142
- support for Windows 2000 195
- support for Windows NT 180
- support for Windows Server 2003 210
- supported environments
 - Novell host systems 133
 - Sun 141
- synchronizing logical volumes 84
- System Management Interface Tool (SMIT) 19
 - definition 19
 - using for configuring 24
 - using to access the Add a Data Path Volume to a Volume Group panel on AIX host 74
 - using to access the Add a Volume Group with Data Path Devices panel on AIX host 73
 - using to access the Back Up a Volume Group with Data Path Devices on AIX host 75
 - using to access the Configure a Defined Data Path Device panel on AIX host 72, 73
 - using to access the Define and Configure All Data Path Devices panel on AIX host 72

- System Management Interface Tool (SMIT) (*continued*)
 - using to access the Display Data Path Device Adapter Status panel on AIX host 72
 - using to access the Display Data Path Device Configuration panel on AIX host 70
 - using to access the Display Data Path Device Status panel on AIX host 71
 - using to access the Remake a Volume Group with Data Path Devices on AIX host 75
 - using to access the Remove a Data Path Device panel on AIX host 73
 - using to access the Remove a Physical Volume from a Volume Group panel on AIX host 74
 - using to back up a volume group with SDD on AIX host 69, 75
 - using to create a volume group with SDD on AIX host 61
 - using to display the ESS vpath device configuration on AIX host 60
 - using to display the SAN Volume Controller vpath device configuration on AIX host 60
 - using to export a volume group with SDD on AIX host 63
 - using to extend an existing SDD volume group on AIX host 68
 - using to import a volume group with SDD on AIX host 62
 - using to remove SDD from AIX host 33
 - using to restore a volume group with SDD on AIX host 69, 75
 - using to unconfigure SDD devices on AIX host 25
 - using to verify SDD configuration on AIX host 30

T

- Taiwan class A compliance statement 233
- trademarks 230

U

- umount
 - /cdrom command 150
 - command 164
- unconfiguring a SDD device to Defined state 68
- unconfiguring all SDD devices to Defined state 68
- unconfiguring SDD on AIX 25
- understanding
 - how SDD works for HP host systems 90
 - how SDD works for Sun host systems 142
- unmod ./sdd-mod.o command 127
- unsupported environments
 - AIX 8
 - HP 89, 90
 - Linux 117
 - Novell host systems 134
 - Sun 142
 - Windows 2000 183
 - Windows NT 165
 - Windows Server 2003 199
- Updating SDD using a PTF 29

- upgrading
 - AIX
 - adapter firmware level 14
 - SDD
 - HP host system 96
 - Linux host system 119
 - Sun host system 150
 - Windows 2000 host system 188
 - Windows NT host system 174
 - Windows Server 2003 host system 204
- upgrading automatically
 - SDD 26
- upgrading to SDD 1.3.3.3 (or later) in a two-node cluster environment 188
- using
 - HP applications with SDD 107
 - Linux
 - single command for SDD configuration 125
 - standard UNIX applications 131
 - Sun applications with SDD 157
- using command 217
- using ESS devices directly, AIX 81
- using ESS devices through AIX LVM 82
- using SAN Volume Controller devices through AIX LVM 82
- using the datapath commands 217
- using the trace function, AIX 86
- utility programs, AIX
 - addpaths 76
 - dpovgfix 77
 - extendvg4vp 78
 - hd2vp and vp2hd 76
 - lsvpcfg 77
 - mkvg4vp 78
 - using ESS devices directly 81
 - using ESS devices through AIX LVM 82
 - using SAN Volume Controller devices through AIX LVM 82
 - using the trace function 86
- utility programs, HP
 - hd2vp 98
 - vp2hd 98

V

- varyoffvg command 24
- verifying
 - additional paths are installed correctly
 - Windows 2000 host system 192
 - Windows NT host system 171
 - Windows Server 2003 host system 208
 - AIX
 - configuring SDD 30
 - SDD installation 20
 - SDD installation
 - Linux host systems 120
- vgexport command
 - HP host system 105, 106
- vgimport command
 - HP host system 106

- volume group
 - mixed
 - how to fix problem 66
 - mixed volume groups
 - dpovgfix vg-name 66
- volume groups on AIX 62
- vp2hd command
 - HP host system 98

W

- Web site
 - AIX APARs, maintenance level fixes and microcode updates 9
 - information about removing multipath access to your shared volume
 - Multiple-Path Software May Cause Disk Signature to Change (Knowledge Base Article Number Q293778) 177
 - information on the Fibre Channel adapters that can be used on your AIX host 9
 - information on the Fibre Channel adapters that can be used on your Novell host 134
 - information on the SCSI adapters that can attach to your AIX host 9
 - information on the SCSI adapters that can attach to your Novell host 134
 - Novell APARs, maintenance level fixes and microcode updates 133
 - SDD xvi
- Web sites
 - HP documentation 109, 112
 - information about
 - SCSI adapters that can attach to your Windows 2000 host system 184
 - SCSI adapters that can attach to your Windows NT host system 166
 - SCSI adapters that can attach to your Windows Server 2003 host system 200
- Windows 2000
 - unsupported environments 183
- Windows 2000 host system
 - adding
 - paths to SDD devices 190
 - clustering special considerations 195
 - configuring
 - cluster with SDD 196
 - ESS 185
 - Fibre Channel adapters 185
 - SCSI adapters 185
 - disk driver 1
 - displaying the current version of the SDD 188
 - installing SDD 1.4.0.0 (or later) 186
 - path reclamation 195
 - protocol stack 1
 - removing SDD 188
 - SDD 183
 - support for clustering 195
 - upgrading SDD 188
 - verifying
 - additional paths to SDD devices 192

- Windows 2000 host systems
 - error recovery and retry policy 194
- Windows NT
 - adding
 - paths to SDD devices 168
 - unsupported environments 165
- Windows NT host system
 - adding
 - multipath storage configuration to the ESS 174
 - new storage to existing configuration 175
 - clustering special considerations 180
 - configuring
 - additional paths 170
 - clusters with SDD 180
 - SDD 168
 - disk driver 1
 - displaying the current version of the SDD 178
 - installing
 - additional paths 170
 - SDD 167
 - modifying multipath storage configuration to the ESS 174
 - path reclamation 180
 - protocol stack 1
 - removing SDD 176
 - reviewing existing SDD configuration information 168, 174
 - SDD 165
 - support for clustering 180
 - upgrading
 - SDD 174
 - verifying
 - additional paths to SDD devices 171
 - new storage is installed correctly 175
- Windows NT host systems
 - error recovery and retry policy 178
- Windows Server 2003
 - unsupported environments 199
- Windows Server 2003 host system
 - adding
 - paths to SDD devices 206
 - clustering special considerations 211
 - configuring
 - cluster with SDD 211
 - ESS 201
 - Fibre Channel adapters 201
 - SCSI adapters 201
 - displaying the current version of the SDD 203
 - installing SDD 1.4.0.0 201
 - path reclamation 211
 - removing SDD 204
 - SDD 199
 - support for clustering 210
 - upgrading SDD 204
 - verifying
 - additional paths to SDD devices 208
- Windows Server 2003 host systems
 - error recovery and retry policy 209

Readers' Comments — We'd Like to Hear from You

**Subsystem Device Driver User's Guide
for the IBM TotalStorage Enterprise Storage Server and the IBM TotalStorage SAN Volume Controller**

Publication No. SC26-7540-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



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
RCF Processing Department
G26/050
5600 Cottle Road
San Jose, CA 95193-0001



Fold and Tape

Please do not staple

Fold and Tape



Printed in U.S.A.

SC26-7540-00



Spine information:



Subsystem Device Driver User's Guide