IBM TotalStorage™ Enterprise Storage Server™:

Subsystem Device Driver User's Guide

Read Before Using

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

SC26-7478-00

IBM TotalStorage™ Enterprise Storage Server™:

Subsystem Device Driver User's Guide

SC26-7478-00

Note
Before using this information and the product it supports, read the information in “Notices” on page 165.

First Edition (August 2002)

This edition replaces the Twelfth Edition of this book (form number GC26-7442-03), and it includes information that specifically applies to IBM ESS Subsystem Device Driver (SDD):

- Version 1 Release 3 Modification 2 Level 2 for AIX 4.3.2, AIX 4.3.3, AIX 5.1.0
- Version 1 Release 3 Modification 1 Level 5 for HP-UX 11.0 and HP-UX11i
- Version 1 Release 3 Modification 0 Level 1 for Linux RedHat 7.2 and SuSE 7.3 (Intel i686 only)
- Version 1 Release 3 Modification 1 Level 0 for Solaris 2.6, Solaris 7, Solaris 8
- Version 1 Release 3 Modification 2 Level 3 for Windows 2000 Service Pack 2 or higher
- Version 1 Release 3 Modification 1 Level 0 for Windows NT 4.0 Service Pack 3 or higher

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

© Copyright International Business Machines Corporation 1999, 2002. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures ix

Tables xi

About this book xiii
Who should use this book xiii
Summary of changes xiii
New information xiii
Modified information xiv
Related information xv
The IBM TotalStorage ESS library xv
Other IBM publications xvi
Ordering IBM publications xx
Non-IBM publications. xx
How to send your comments. xxi

Chapter 1. Overview of SDD 1
The SDD architecture 1
Enhanced data availability 3
Dynamic I/O load-balancing. 3
Automatic path-failover protection 3
Concurrent download of ESS licensed internal code. 4

Chapter 2. Using SDD on an AIX host system 5
Command syntax conventions. 5
Highlighting conventions 5
Special characters conventions 5
Verifying the hardware and software requirements 5
Hardware 5
Software. 6
Non-supported environments 6
Host system requirements 6
Preparing for SDD installation 7
Configuring the ESS 7
Installing the AIX fibre-channel device drivers 8
Uninstalling the AIX fibre-channel device drivers 8
Configuring fibre-channel attached devices 9
Removing fibre-channel attached devices 9
Verifying the Emulex adapter firmware level. 9
Upgrading the Emulex adapter firmware level 10
| Determining if the SDD server for Expert is installed 10
| Planning for SDD installation on a pSeries 690 LPAR. 11
Determining the installation package 11
Installing SDD 14
Verifying your currently installed version of SDD. 15
Preparing to configure SDD 17
Configuring SDD 17
Unconfiguring Subsystem Device Driver. 17
Verifying the SDD configuration 18
Changing the path-selection policy. 19
Adding paths to SDD devices of a volume group 20
| Opening an invalid or close_dead path dynamically 21
Upgrading SDD for AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0 21

© Copyright IBM Corp. 1999, 2002

Understanding SDD support for installation upgrade 21
Verifying your previously installed version of SDD 22
Upgrading SDD. 23
Upgrading SDD manually 23
Removing SDD from an AIX host system 24
Using ESS concurrent download of licensed internal code 25
SDD server daemon 26
Understanding how the SDD server daemon works 26
Verifying if the SDD server has started 26
Starting the SDD server manually 27
Changing to a different port number for the SDD server 27
Stopping the SDD server 27
Understanding SDD 1.3.2.6 (or later) support for single-path configuration . . . 27
Understanding SDD error recovery policies 28
Understanding the SDD support for pSeries 690 with static LPARs configured 29
Understanding the SDD support for High Availability Cluster Multi-Processing (HACMP) 30
SDD installation package attributes 34
HACMP RAID concurrent mode volume groups and enhanced concurrent capable volume groups 37
HACMP RAID concurrent mode volume groups 37
Enhanced concurrent capable volume groups 41
Recovering paths that are lost during HACMP node failover 43
Supporting enhanced concurrent mode in an HACMP environment. 43
Managing secondary system paging space 44
Listing paging spaces 44
Adding a paging space 44
Removing a paging space. 45
Providing load-balancing and failover protection. 45
Displaying the ESS vpath device configuration 45
Configuring volume groups for failover protection 46
Importing volume groups with SDD 48
Importing volume groups under HACMP 48
Exporting a volume group with SDD 49
Losing failover protection 49
Persistent reserve command tool 61
Using ESS devices directly 63
Using ESS devices through AIX LVM. 64
Migrating a non-SDD volume group to an ESS SDD multipath volume group in concurrent mode 65
Using the trace function 69
Error log messages 69

Chapter 3. SDD for a Linux host system 71
Verifying the hardware and software requirements 71
Hardware 71
Software 71
Non-supported environments. 71
Preparing for SDD installation 71
Configuring the ESS 72
Configuring fibre-channel adapters. 72
Installing SDD 72
Verifying the SDD installation. 73
Configuring SDD 73
Using multiple commands for SDD configuration 74
Using a single command for SDD configuration 76

Configuring SDD at system startup . . . 76
Maintaining vpath device configuration persistence . . . 77
Removing SDD . . . 77
Understanding the SDD error recovery policy . . . 78
Understanding the SDD support for single-path configuration . . . 78
Partitioning SDD vpath devices . . . 78
Using standard UNIX applications . . . 79

Chapter 4. SDD for a Windows NT host system . . . 81
Verifying the hardware and software requirements . . . 81
Hardware . . . 81
Software . . . 81
Non-support environments . . . 81
ESS requirements . . . 81
Host system requirements . . . 81
Preparing for SDD installation . . . 82
Configuring the ESS . . . 82
Configuring fibre-channel adapters . . . 82
Configuring SCSI adapters . . . 83
Installing the SDD . . . 83
Configuring the SDD . . . 84
Adding paths to SDD devices . . . 84
Upgrading the SDD . . . 88
Adding or modifying a multipath storage configuration to the ESS . . . 88
Removing the Subsystem Device Driver . . . 92
Displaying the current version of the Subsystem Device Driver . . . 93
Error recovery and retry policy . . . 93
Using high availability clustering . . . 94
Special considerations in the high availability clustering environment . . . 94
Configuring a Windows NT cluster with SDD . . . 95

Chapter 5. SDD for a Windows 2000 host system . . . 97
Verifying the hardware and software requirements . . . 97
Non-supported environments . . . 97
ESS requirements . . . 97
Host system requirements . . . 97
Preparing for SDD 1.3.3.3 (or later) installation . . . 98
Configuring the ESS . . . 98
Configuring fibre-channel adapters . . . 99
Configuring SCSI adapters . . . 99
Planning for SDD 1.3.3.3 (or later) installation . . . 99
Installing SDD 1.3.3.3 (or later) . . . 100
Displaying the current version of the SDD . . . 101
Upgrading the Subsystem Device Driver . . . 101
Upgrading to SDD 1.3.3.3 (or later) in a 2–node cluster environment . . . 102
Uninstalling the Subsystem Device Driver . . . 103
Removing the Subsystem Device Driver in a 2–node cluster environment 103
Adding paths to SDD devices . . . 103
Activating additional paths . . . 104
Verifying additional paths are installed correctly . . . 105
Error recovery and retry policy . . . 106
Support for Windows 2000 clustering . . . 107
Special considerations in the Windows 2000 clustering environment . . . 107
Configuring a Windows 2000 cluster with SDD . . . 108

Chapter 6. SDD for a Hewlett Packard host system . . . 111

Verifying the hardware and software requirements . . . 111
Hardware . . . 111
Software . . . 111
Non-supported environments . . . 111
SDD support for 32-bit and 64-bit applications on HP-UX 11.0 and HP-UX 11i (Fibre) . . . 111
Understanding how SDD works on an HP host system . . . 112
Preparing for SDD installation . . . 112
Configuring the ESS . . . 112
Planning for installation . . . 112
Determining if the SDD 1.3.1.5 (or later) server for Expert is installed . . . 113
Installing the SDD . . . 114
Post-installation . . . 116
Upgrading the SDD . . . 117
Uninstalling the Subsystem Device Driver . . . 117
Changing an SDD hardware configuration . . . 118
SDD server daemon . . . 119
Understanding how the SDD server daemon works . . . 119
Verifying if the SDD server has started . . . 119
Starting the SDD server manually . . . 119
Changing to a different port number for the SDD server . . . 120
Stopping the SDD server . . . 120
Understanding the SDD 1.3.1.5 support for single-path configuration . . . 120
Understanding the SDD error recovery policy . . . 121
Converting a volume group . . . 121
How to import and export volume groups . . . 122
Exporting volume groups . . . 122
Moving the map file . . . 123
Creating the volume group device directory . . . 123
Creating the group special file . . . 123
Importing volume groups . . . 123
Using applications with SDD . . . 124
Standard UNIX applications . . . 124
Installing SDD on a Network File System file server . . . 129
Oracle . . . 130

Chapter 7. SDD for a Sun host system . . . 135
Verifying the hardware and software requirements . . . 135
Hardware . . . 135
Software . . . 135
Supported environments . . . 135
Nonsupported environments . . . 135
Understanding how SDD works on a Sun host system . . . 136
Preparing for SDD installation . . . 136
Configuring the ESS . . . 136
Determining if the SDD server for Expert is installed . . . 136
Planning for installation . . . 137
Installing the SDD . . . 138
Verifying the SDD installation . . . 139
Uninstalling an SDD . . . 139
Post-installation . . . 140
Upgrading the SDD . . . 141
Changing an SDD hardware configuration . . . 142
SDD server daemon . . . 142
Understanding how the SDD server daemon works . . . 142
Verifying if the SDD server has started . . . 143

Starting the SDD server manually 143
Changing to a different port number for the SDD server 143
Stopping the SDD server 144
Understanding the SDD 1.3.1.0 (or later) support for single-path configuration 144
Understanding the SDD 1.3.1.0 error recovery policy 144
Using applications with SDD 145
Standard UNIX applications 145
Installing SDD on a Network File System file server 145
Oracle 146
Solstice DiskSuite 150

Chapter 8. Using the datapath commands 153
datapath query adapter command 154
datapath query adaptstats command 155
datapath query device command 156
datapath query devstats command 158
datapath set adapter command 160
datapath set path command 161
datapath open path command 162

Notices 165
Trademarks. 165
Electronic emission notices 166
Federal Communications Commission (FCC) statement 166
Industry Canada compliance statement 167
European community compliance statement 167
Japanese Voluntary Control Council for Interference (VCCI) class A statement 168
Korean government Ministry of Communication (MOC) statement 168
Taiwan class A compliance statement 168
IBM agreement for licensed internal code. 168
Actions you must not take 169

Glossary 171

Index 191

Figures

1. Multipath connections between a host system and the disk storage in an ESS 3
2. Output from the Display Data Path Device Configuration SMIT panel. 45
3. IBMdpo Driver 32-bit 115
4. IBMdpo Driver 64-bit 115

© Copyright IBM Corp. 1999, 2002

Tables

1. Publications in the ESS library xv
2. Other IBM publications related to the ESS xvii
3. Non-IBM publications related to ESS xxi
4. SDD in the protocol stack 1
5. AIX required fixes 6
6. SDD 1.3.2.6 support for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0 11
7. SDD 1.3.2.6 (or later) installation packages for non–HACMP and HACMP environments on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0 12
8. Major files included in the SDD installation package 13
9. List of previously installed installation packages that are supported with the installation upgrade 22
10. Software support for HACMP in concurrent mode 31
11. Software support for HACMP in nonconcurrent mode 32
12. Software support for HACMP in concurrent mode on AIX 5.1.0 (32-bit or 64-bit kernel) 32
13. Software support for HACMP in nonconcurrent mode on AIX 5.1.0 (32-bit and 64-bit kernel) 33
14. HACMP and supported SDD features 34
15. HACMP and supported features 36
16. HACMP and supported features 36
17. SDD-specific SMIT panels and how to proceed. 54
18. SDD components for a Linux host system 73
19. Summary of SDD commands for a Linux host system 74
20. SDD 1.3.3.3 installation scenarios 99
21. SDD installation scenarios 112
22. Patches necessary for proper operation of SDD on HP–UX 11.0 113
23. SDD components installed for HP host systems 116
24. System files updated for HP host systems 116
25. SDD commands and their descriptions for HP host systems 116
26. SDD installation scenarios 137
27. SDD package file names 137
28. SDD components installed for Sun host systems 141
29. System files updated for Sun host systems. 141
30. SDD commands and their descriptions for Sun host systems 141
31. Commands 153

© Copyright IBM Corp. 1999, 2002

About this book

This book provides step-by-step procedures on how to install, configure, and use the IBM® TotalStorage™ Enterprise Storage Server™ (ESS) Subsystem Device Driver (SDD) for the following host systems:

- IBM AIX®
- HP
- Linux Red Hat and SuSE (Intel)
- Sun
- Microsoft® Windows NT®
- Microsoft Windows 2000

Who should use this book

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

Summary of changes

This book contains both information previously presented in the Twelfth Edition (May 2002) of the *IBM TotalStorage™ Enterprise Storage Server™: Subsystem Device Driver User's Guide* and major technical changes to that information. Technical changes are indicated by revision bars (|) in the left margin of the book.

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:

What's new in Chapter 2. SDD for an AIX host system:

- A section about "Determining if the SDD server for Expert is installed" on page 10
- A reference to the **datapath open path** command in "Opening an invalid or close_dead path dynamically" on page 21
- The reference to lquerypr in Table 8 on page 13 and other appropriate sections
- The reference to sddsrv in Table 8 on page 13 and other appropriate sections
- A section about "SDD server daemon" on page 26
- A section about "Understanding SDD 1.3.2.6 (or later) support for single-path configuration" on page 27
- A section about "Understanding SDD error recovery policies" on page 28
- A section about "Planning for SDD installation on a pSeries 690 LPAR" on page 11
- A section about "Managing secondary system paging space" on page 44

What's new in Chapter 3. SDD for a Linux host system:

- This chapter is new.

What's new in Chapter 6. SDD for a Hewlett-Packard host system:

© Copyright IBM Corp. 1999, 2002

- A section about “Determining if the SDD 1.3.1.5 (or later) server for Expert is installed” on page 113
- A section about “SDD server daemon” on page 119
- A section about “Understanding the SDD 1.3.1.5 support for single-path configuration” on page 120
- A section about “Understanding the SDD error recovery policy” on page 121
- A section about “How to import and export volume groups” on page 122
- A section about “Converting a volume group” on page 121

What’s new in Chapter 7. SDD for a Sun host system:

- A section about “Determining if the SDD server for Expert is installed” on page 136
- A section about “SDD server daemon” on page 142
- A section about “Understanding the SDD 1.3.1.0 (or later) support for single-path configuration” on page 144
- A section about “Understanding the SDD 1.3.1.0 error recovery policy” on page 144

What’s new in Chapter 7. Using the datapath commands:

- A section about “datapath open path command” on page 162

Modified information

This edition includes the following modified information:

What’s modified in Chapter 1. Overview of the Subsystem Device Driver:

What’s modified in Chapter 2. SDD for an AIX host system:

- Updated descriptions for ibmSdd_510.rte and ibmSdd_510nchacmp.rte installation packages in Table 7 on page 12.
- Removed references to single-path concurrent download of licensed internal code from appropriate sections; ESS no longer supports single-path concurrent download of licensed internal code.
- Removed references to AIX 4.2.1 installation package (ibm_Sdd_421.rte) from appropriate sections; SDD no longer supports AIX 4.2.1.
- Replaced **odmget –q ″name –dpo″ CuAt** command with **lsattr -El dpo** command in “SDD installation package attributes” on page 34 section.

What’s modified in Chapter 6. SDD for a Hewlett-Packard host system

- Corrected installation requirement for patches that are listed Table 22 on page 113. (These patches must be installed on HP 11.0 host system to ensure that SDD operates successfully. Do not install these patches on HP-UX11i host systems.) This is a correction to the Eleventh Edition (April 2002) of the ″Subsystem Device Driver User’s Guide″ (form number GC26-7442-02).
- Removed references to single-path concurrent download of licensed internal code from appropriate sections; ESS no longer supports single-path concurrent download of licensed internal code.

What’s modified in Chapter Chapter 7. SDD for a Sun host system

- Removed references to single-path concurrent download of licensed internal code from appropriate sections; ESS no longer supports single-path concurrent download of licensed internal code.

Related information

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

- The publications that make up the IBM TotalStorage ESS library
- Other IBM publications that relate to the ESS
- Non-IBM publications that relate to the ESS

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

The IBM TotalStorage ESS library

Table 1 lists and describes the customer publications that 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-8770. These publications are also available on the following ESS Web site as PDF files:

http://ssddom02.storage.ibm.com/disk/ess/documentation.html

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

Table 1. Publications in the ESS library

Title	Description	Order number
IBM TotalStorage Enterprise Storage Server: Copy Services Command-Line Interface Reference	This guide describes the commands that you can use from the ESS Copy Services command-line interface (CLI). 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.	SC26-7449 (See Note.)
IBM TotalStorage Enterprise Storage Server: Configuration Planner for Open-Systems Hosts	This guide provides guidelines and work sheets for planning the logical configuration of an ESS that attaches to open-system hosts.	SC26-7475 (See Note.)
IBM TotalStorage Enterprise Storage Server: Configuration Planner for S/390® and IBM @server zSeries Hosts	This guide provides guidelines and work sheets for planning the logical configuration of an ESS that attaches to either the IBM S/390® and IBM @server zSeries™ host system.	SC26-7476 (See Note.)
IBM TotalStorage Enterprise Storage Server: Host Systems Attachment Guide	This guide provides guidelines for attaching the ESS to your host system and for migrating to fibre-channel attachment from either a small computer system interface (SCSI) or from the IBM SAN Data Gateway.	SC26-7446
IBM TotalStorage Enterprise Storage Server: DFSMS Software Support Reference	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

Table 1. Publications in the ESS library (continued)

Title	Description	Order number
IBM TotalStorage™ Enterprise Storage Server™: Introduction and Planning Guide	This guide introduces the ESS product and lists the features you can order. It also provides guidelines for planning the installation and configuration of the ESS.	GC26-7444
IBM TotalStorage Enterprise Storage Server: Quick Configuration Guide	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
IBM TotalStorage Enterprise Storage Server: System/390® Command Reference	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
IBM TotalStorage Storage Solutions Safety Notices	This publication provides translations of the danger notices and caution notices that IBM uses in ESS publications.	GC26-7229
IBM TotalStorage Enterprise Storage Server: SCSI Command Reference	This publication describes the functions of the ESS. It provides reference information, such as channel commands, sense bytes, and error recovery procedures for UNIX®, IBM Application System/400® (AS/400®), and @server iSeries™ 400 hosts.	SC26-7297
IBM TotalStorage Enterprise Storage Server: Subsystem Device Driver User's Guide	This publication describes how to use the IBM TotalStorage ESS Subsystem Device Driver (SDD) on open-systems hosts to enhance performance and availability on the ESS. SDD creates redundant paths for shared logical unit numbers. SDD permits applications to run without interruption when path errors occur. It balances the workload across paths, and it transparently integrates with applications. For information about the SDD, go to the following Web site: http://ssddom01.storage.ibm.com/techsup/swtechsup.nsf/support/sddupdates/	SC26-7478
IBM TotalStorage Enterprise Storage Server: User's Guide	This guide provides instructions for setting up and operating the ESS and for analyzing problems.	SC26-7445
IBM TotalStorage Enterprise Storage Server: Web Interface User's Guide	This guide provides instructions for using the two ESS Web interfaces, ESS Specialist and ESS Copy Services.	SC26-7448
Note: No hardcopy book is produced for this publication. However, a PDF file is available from the following Web site: www.storage.ibm.com/hardsoft/products/ess/refinfo.htm		

Other IBM publications

Table 2 on page xvii lists and describes other IBM publications that contain additional information that is related to the ESS. See "The IBM publications center" on page xx for information about ordering these and other IBM publications.

Note: Informal documents, such as white papers, are not available from the "The IBM publications center" on page xx.

Table 2. Other IBM publications related to the ESS

Title	Description	Order number
DFSMS/MVS® Version 1: Advanced Copy Services	This publication helps you understand and use IBM Advanced Copy Services functions on an S/390 or zSeries. It describes two dynamic copy functions and several point-in-time copy functions. These functions provide backup and recovery of data if a disaster occurs to your data center. The dynamic copy functions are Peer-to-Peer Remote Copy and Extended Remote Copy. Collectively, these functions are known as remote copy. FlashCopy and Concurrent Copy are the point-in-time copy functions.	SC35-0355
DFSMS/MVS V1: Remote Copy Guide and Reference	This publication provides guidelines for using remote copy functions with S/390 and zSeries hosts.	SC35-0169
Enterprise Storage Solutions Handbook	This handbook helps you understand what makes up enterprise storage management. The concepts include the key technologies that you need to 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
ES/9000, ES/3090: IOCP User Guide Volume A04	This document describes the Input/Output Configuration Program (IOCP) 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
ESA/390: ESCON I/O Interface	This publication provides a description of the physical and logical ESA/390 I/O interface and the protocols that govern information transfer over that interface. It is intended for designers of programs and equipment associated with the ESCON® I/O interface and for service personnel who maintain that equipment. However, anyone concerned with the functional details of the ESCON I/O interface will find it useful.	SA22-7202
ESS Fibre-Channel Migration Scenarios	This white paper describes how to change your host system attachment to the ESS from SCSI and SAN Data Gateway to native fibre-channel attachment. To get the white paper, go to the following Web site: www.storage.ibm.com/hardsoft/products/ess/refinfo.htm	No order number
ESS Solutions for Open Systems Storage: Compaq AlphaServer, HP, and Sun	This publication 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 running the Open VMS operating system. This book also focuses on the settings required to give optimal performance and on the settings for device driver levels. This book is for the experienced UNIX professional who has a broad understanding of storage concepts.	SG24-6119
Fibre Channel Connection (FICON) I/O Interface: Physical Layer	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
Fibre-channel Subsystem Installation Guide	This publication tells you how to attach the IBM xSeries™ 430 and NUMA-Q® host systems with fibre-channel adapters to the ESS. Contact your sales representative to obtain this publication.	No order number
Fibre Transport Services (FTS): Physical and Configuration Planning Guide	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234

Table 2. Other IBM publications related to the ESS (continued)

Title	Description	Order number
IBM Enterprise Storage Server	This publication, from the IBM International Technical Support Organization, introduces the ESS and provides an understanding of its benefits. It also describes in detail the architecture, hardware, and functions of the ESS.	SG24-5465
IBM Enterprise Storage Server Performance Monitoring and Tuning Guide	This guide, from the IBM International Technical Support Organization, provides guidance on the best way to configure, monitor, and manage your ESS to ensure optimum performance.	SG24-5656
IBM SAN Fibre Channel Managed Hub: 3534 Service Guide	This guide explains how to convert the IBM SAN Fibre Channel Managed Hub from a Fibre-Channel Arbitrated Loop (FC-AL) configuration to a fabric-capable switched environment.	GC26-7391
IBM SAN Fibre Channel Managed Hub: 3534 User's Guide	The IBM SAN Fibre Channel Switch 3534 is an eight-port Fibre Channel Gigabit Hub that consists of a motherboard with connectors for supporting up to eight ports, including seven fixed shortwave optic ports and one GBIC port, and an operating system for building and managing a switched loop architecture.	SY27-7616
IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08. To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/	SC26-7350
IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide	This guide describes the IBM SAN Fibre Channel Switch and the IBM StorageWatch Specialist. It provides information about the commands and how to manage the switch with Telnet and the Simple Network Management Protocol (SNMP). To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/	SC26-7349
IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide	This publication describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S16. It is intended for trained service representatives and service providers who act as the primary level of field hardware service support to help solve and diagnose hardware problems. To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/	SC26-7352
IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide	This guide introduces the IBM SAN Fibre Channel Switch 2109 Model S16 and tells you how to manage and monitor the switch using zoning and how to manage the switch remotely. To get a copy of this guide, go to the following Web site: www.ibm.com/storage/fcswitch/	SC26-7351
IBM StorWatch Expert: Hands-On Usage Guide	This guide, from the IBM Technical Support Organization, helps you install, tailor, configure, and use StorWatch™ Expert.	SG24-6102
IBM TotalStorage Expert Version 2 Release 1: Installation Guide	This guide helps you install the IBM TotalStorage Expert (formally the IBM StorWatch Expert) program. This IBM TotalStorage Expert provides asset, capacity, and performance management information for disk and tape storage systems.	GC26-7436
IBM Versatile Storage Server: Introduction and Planning Guide	This publication introduces the IBM Versatile Storage Server™ and lists the features you can order. It also provides planning information for both 2105 Models B09 and 100.	GC26-7223

Table 2. Other IBM publications related to the ESS (continued)

Title	Description	Order number
Implementing ESS Copy Services on S/390	This publication 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
Implementing ESS Copy Services on UNIX and Windows NT/2000	This publication tells you how to install, customize, and configure ESS Copy Services on UNIX or Windows NT® host systems. The Copy Services functions that are supported include Peer-to-Peer Remote Copy, FlashCopy, Extended Remote Copy, and Concurrent Copy. 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 (HACMP) environment.	SG24-5757
Implementing Fibre Channel Attachment on the ESS	This publication, from the IBM 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 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
Implementing the IBM Enterprise Storage Server	This publication can help you install, tailor, and configure the ESS in your environment.	SG24-5420
IOCP User's Guide, IBM @server zSeries	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
OS/390: Hardware Configuration Definition User's Guide	This guide provides detailed information about the input/output definition file (IODF) and about how to configure parallel access volumes (PAVs). This guide discusses how to use Hardware Configuration Definition (HCD) for both OS/390® and z/OS V1R1.	SG24-6102
OS/390: MVS System Messages Volume 1 (ABA - ASA)	This publication lists OS/390 MVS™ system messages ABA to ASA.	GC28-1784
S/390: IOCP User's Guide and ESCON CTC Reference	This publication describes the Input/Output Configuration Program that supports ESCON architecture and the ESCON multiple image facility.	GC38-0401
S/390: Stand-Alone Input/Output Configuration Program User's Guide	This guide helps you use the stand-alone IOCP on an IBM S/390 9672 Parallel Transaction Server, an IBM S/390 9672 Parallel Enterprise Server, and an IBM S/390 9674 Coupling Facility.	GC38-0455
S/390: Stand-Alone Input/Output Configuration Program User's Guide	This guide helps you use the stand-alone IOCP on an IBM S/390 Parallel Enterprise Server - Generation 3 or Generation 4, an IBM S/390 Coupling Facility Model C04 or Model C05, and on an IBM S/390 Application StarterPak 3000.	GC38-0456

Table 2. Other IBM publications related to the ESS (continued)

Title	Description	Order number
S/390: Stand-Alone Input/Output Configuration Program User's Guide	This guide helps you use the stand-alone IOCP on an IBM S/390 Parallel Enterprise Server - Generation 5 or Generation 6. All of the following terms are used in this guide to mean the IBM S/390 Parallel Enterprise Server - Generation 5 or Generation 6: • Processor • System • All models • Rack-mounted system	GC38-0458
z/Architecture Principles of Operation	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

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.

The 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

Table 3 on page xxi lists and describes other related publications that are not available through IBM ordering systems. To order, contact the sales representative at the branch office in your locality.

Table 3. Non-IBM publications related to ESS

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

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this book 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 TotalStorage Enterprise Storage Server (ESS) 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 the IBM ESS. It resides in a host system with the native disk device driver and provides the following functions:

- Enhanced data availability
- Dynamic I/O load-balancing across multiple paths
- Automatic path failover protection
- Concurrent download of licensed internal code

As the diagrams in Table 4 show, SDD resides above the disk driver of a host system in the protocol stack and acts as a pseudo device driver. 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 4. SDD in the protocol stack

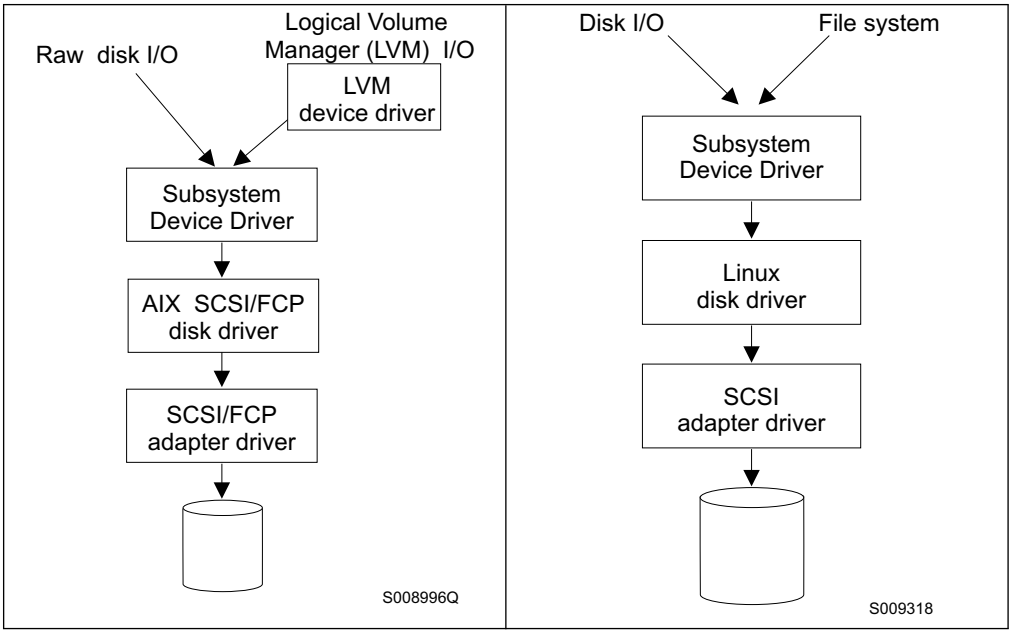

© Copyright IBM Corp. 1999, 2002

Table 4. SDD in the protocol stack (continued)

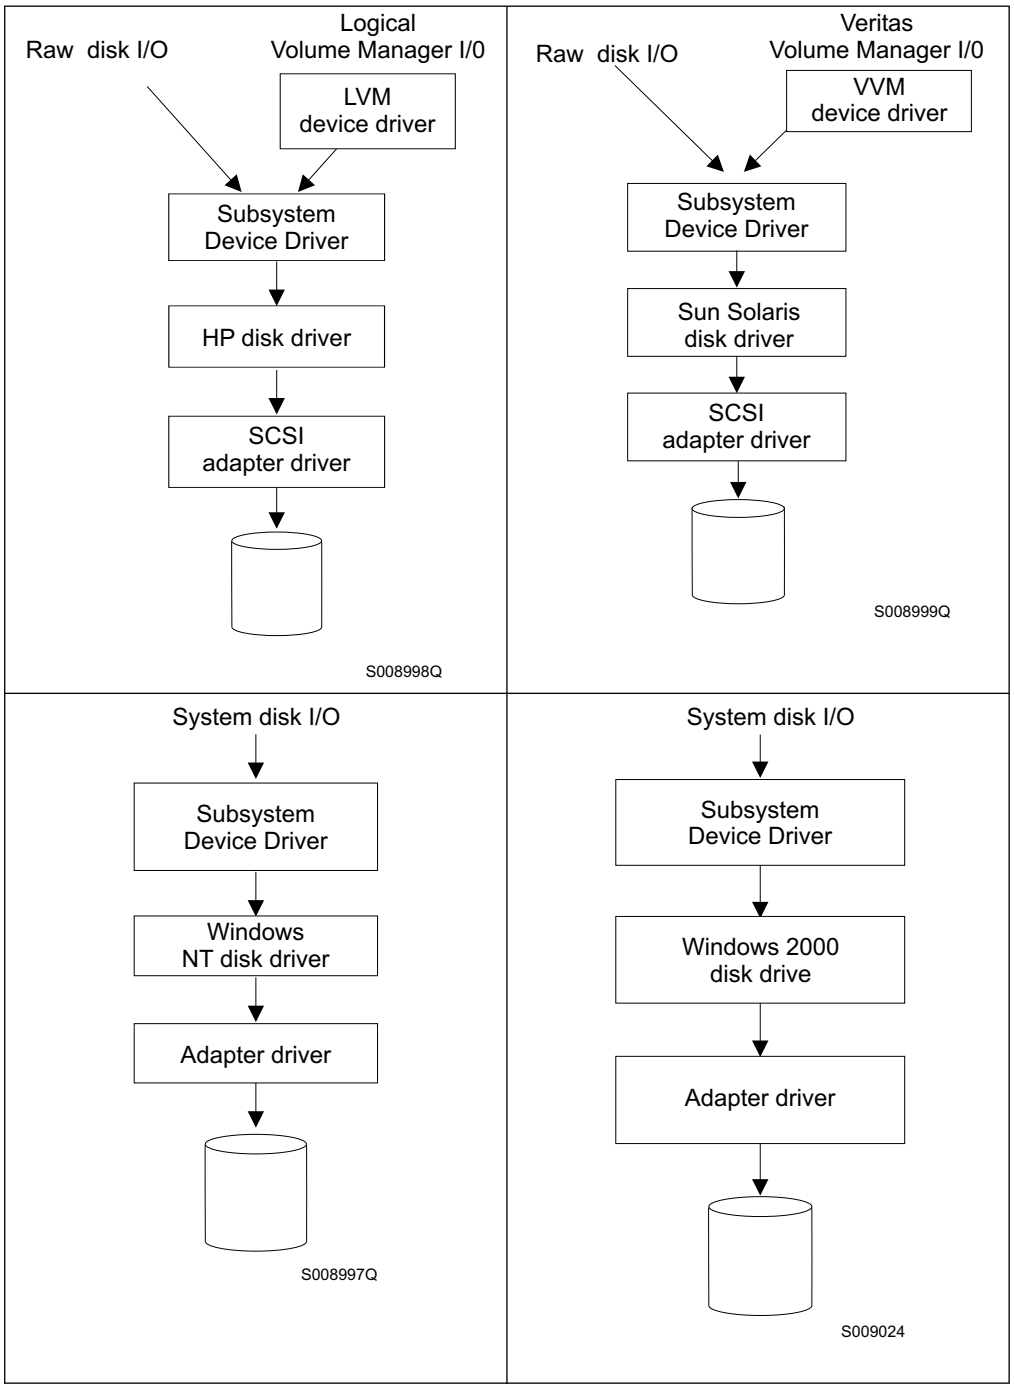

Each SDD device represents a unique physical device on the storage server. There can be up to 32 operating system disk devices that represent up to 32 different paths to the same physical device.

SDD devices behave almost like native operating system disk devices. Most operations on an operating system disk device can be performed on the SDD device, including commands such as **open**, **close**, **dd**, or **fsck**.

Commands that are used in the following situation are not supported for the SDD devices:

- Commands getting the information directly from the system file where the device driver is not involved for the operation.

Enhanced data availability

Figure 1 shows that an SDD-residing host system is attached through SCSI or fibre-channel adapters to an ESS that has internal component redundancy and multipath configuration. SDD 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.

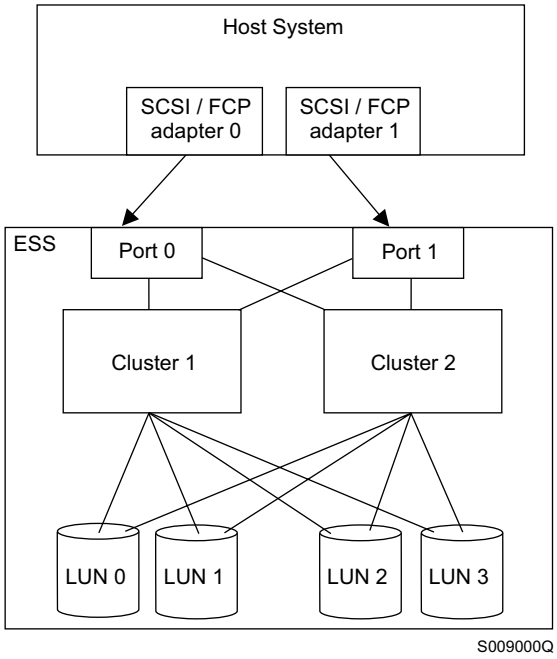

Figure 1. Multipath connections between a host system and the disk storage in an ESS

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 is designed to minimize any disruptions in I/O operations and recover I/O operations from a failing data path. SDD provides path-failover protection through the following process:

- Detecting a path failure
- Notifying the host system of the path failure
- Selecting and using an alternate data path

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

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 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.

Chapter 2. Using SDD on an AIX host system

This chapter provides step-by-step procedures on how to install, configure, upgrade, and remove the SDD on an AIX host system that is attached to an ESS. 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

Command syntax conventions

This section describes the notational conventions that this chapter 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 chapter:

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

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

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

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
- Host system
- SCSI adapters and cables
- Fibre adapters and cables

© Copyright IBM Corp. 1999, 2002

Software

- AIX operating system
- SCSI and fibre-channel device drivers
- ibm2105.rte ESS package

Non-supported environments

SDD does not support:

- A host system with both a SCSI and fibre-channel connection to a shared ESS 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 reserve and release device on AIX 4.3.2, 4.3.3 and 5.1.0
- More than 600 vpaths
- Single-path mode during concurrent download and installation of LIC nor during any ESS concurrent maintenance that impacts the path attachment, such as, an ESS host bay adapter replacement.

Host system requirements

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

Tip: You must check for and download the latest APARs, maintenance level fixes, and microcode updates from the following Web site:

www.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)

  Note: *xxx* represents the ESS model number.

SCSI requirements

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

- The bos.adt package is installed. The host system can be a uniprocessor or a multiprocessor system, such as SMP.
- 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 AIX host system, go to the following Web site:

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

Fibre requirements

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

www.techsupport.services.ibm.com/server/support

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

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

www-1.ibm.com/servers/eserver/pseries/library/hardware_docs/

To use the SDD fibre 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 uniprocessor or a multiprocessor system, such as SMP.
- 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 AIX 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 perform the following tasks:

- Configuring the ESS
- Installing the AIX fibre-channel device drivers
- Configuring fibre-channel attached devices
- Verifying the Emulex adapter firmware level
- Determining if the SDD server for Expert is installed
- Determining the installation package

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.

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

Note: Ensure that the ibm2105.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:

www.techsupport.services.ibm.com/server/support

Perform the following steps to install the AIX fibre-channel device drivers:

1. Install the fibre-channel device drivers from the AIX compact disc. The fibre-channel device drivers include the following installation packages:

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

   devices.pci.df1000f7
   The adapter device driver for RS/6000 / pSeries with feature code 6227

   devices.common.IBM.fc
   The FCP and SCSI protocol driver

   devices.fcp.disk
   The FCP disk driver
2. Check to see if the correct APARS are installed by issuing the **instfix -i | grep IY***nnnnn* command (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 9. Otherwise, go to step 3.
3. 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. Use one of the following commands to uninstall 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 displays.
3. Select the filename of the fibre-channel device driver that you want to uninstall. Press Enter. The selected filename displays 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 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.

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

After the system restarts, use the **lsdev -Cc disk** command to check the ESS 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 Emulex 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 **rmdev -dl fcs***N* **-R** command for each installed FCP adapter, 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 -dl fcs0 -R**
- **rmdev -dl fcs1 -R**

Verifying the Emulex 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 support firmware level for fibre-channel adapters, go to the following Web site:

www.techsupport.services.ibm.com/server/mdownload2/download.html

Tip:

- The current firmware level for Emulex LP7000E adapter is sf322A1.
- The current firmware level for Emulex 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 something like this:

   ```
   (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 Emulex adapter firmware level”.

Upgrading the Emulex 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 fibre-channel attached devices (see “Configuring fibre-channel attached devices” on page 9). 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. Locate the file called df1000f7.131.320.320.320.503, which should be in the /etc/microcode directory. 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. 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. Type the filename for the firmware that is contained in the /etc/microcode directory and press Enter or use the **Tab** key to toggle to **Latest**.
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 fcs***N* command to verify the firmware level on each fibre-channel adapter.

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 AIX host system, you must remove this *standalone* version of the SDD server before you proceed with SDD 1.3.2.6. (or later) installation. The installation package for SDD 1.3.2.6 (or later) 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:

```
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
                                 for AIX
Path: /etc/objrepos
     sddServer.rte        1.0.0.0  COMMITTED IBM SDD...
                                 for AIX
```

For instructions on how to remove the *standalone* version of the SDD server (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:

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

For more information about the SDD server daemon, go to "SDD server daemon" on page 26

Planning for SDD installation on a pSeries 690 LPAR

As a standard feature, the IBM @server pSeries™ 690 server supports status 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 operates 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, you need to determine the installation package that is appropriate for your environment. IBM recommends that you use the persistent reserve version of SDD for AIX 5.1.0 (that is, the ibmSdd_510nchacmp.rte installation package) when you:

- Plan to configure multiple LPARs.
- Plan to share your ESS devices in a nonconcurrent mode between multiple LPARs of a pSeries 690 server.

Attention: Failure to use the correct installation package could leave your disks exposed to data corruption, whether your server is running or not running HACMP. Table 7 on page 12 lists the various scenarios and the installation packages that you must use for your environment.

Determining the installation package

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

Table 6 lists and describes SDD support for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0.

Table 7 on page 12 lists the various scenarios and the installation packages that you must use for your environment.

Support for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0 host systems

Table 6 lists and describes SDD support for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0.

Table 6. SDD 1.3.2.6 support for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0

SDD installation package names	AIX kernel mode	Application mode	SDD interface
ibmSdd_432.rte	32-bit	32-bit, 64-bit	LVM, raw device
ibmSdd_433.rte	32-bit	32-bit, 64-bit	LVM, raw device

Table 6. SDD 1.3.2.6 support for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0 (continued)

SDD installation package names	AIX kernel mode	Application mode	SDD interface
ibmSdd_510.rte	32-bit, 64-bit	32-bit, 64-bit	LVM, raw device
ibmSdd_510nchacmp.rte	32-bit, 64-bit	32-bit, 64-bit	LVM, raw device

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

SDD 1.3.2.6 supports AIX 5.1.0 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 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 host system runs in 64-bit mode.

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

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

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

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

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

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

Installation packages for non-HACMP and HACMP environments on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0 host systems

Attention:

- Use the persistent reserve (PR) versions of SDD for AIX 4.3.2 (ibmSdd_433.rte) and AIX 5.1.0 (ibmSdd_510nchacmp.rte) when:
  - Your server is sharing disks with other servers or LPARs
- Failure to use the correct installation package could leave disks exposed to data corruption. Table 7 lists the various scenarios and the installation packages that you must use for your environment.

Table 7. SDD 1.3.2.6 (or later) installation packages for non–HACMP and HACMP environments on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0

Host system	Scenario	Environment	Use installation package
AIX 4.3.2	1	Your server is not running HACMP and is not sharing disks with other servers.	ibmSdd_432.rte

Table 7. SDD 1.3.2.6 (or later) installation packages for non–HACMP and HACMP environments on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0 (continued)

Host system	Scenario	Environment	Use installation package
AIX 4.3.3	2	Your server is not running HACMP and is not sharing disks with other servers.	ibmSdd_432.rte
	3	Your server is running HACMP in concurrent mode.	ibmSdd_432.rte
	4	Your server is running HACMP in nonconcurrent mode, or both concurrent and nonconcurrent modes (mixed environment).	ibmSdd_433.rte
	5	Your server is sharing disks with other servers without having HACMP installed. (See Note.)	ibmSdd_433.rte
AIX 5.1.0	6	Your server or LPAR is not running HACMP and is not sharing disks with other servers or LPARs	ibmSdd_510.rte
	7	Your server or LPAR is running HACMP in concurrent mode.	ibmSdd_510.rte
	8	Your server or LPAR is running HACMP in nonconcurrent mode, or both concurrent mode and nonconcurrent modes (mixed environment).	ibmSdd_510nchacmp.rte
	9	Your server or LPAR is sharing disks with other servers or LPARs. (See Note.)	ibmSdd_510nchacmp.rte

Note:

With SDD 1.3.2.6, the **lquerypr** command is available. This command implements several SCSI-3 persistent reserve service actions, such as the following:

- Query the persistent reserve key or reservation status
- Release persistent reserve
- Preempt-abort persistent reserve
- Clear persistent reserve

In Scenario 5 and Scenario 9, since the HACMP is not installed and persistent reserve version of SDD is installed, then the user may encounter a situation where persistent reserve is left on the disk and a standby server or LPAR can not take over the resources. To solve this problem, the user may need the **lquerypr** command to release or clear persistent reserve on the disk. For more detailed information about the **lquerypr** command, see “Persistent reserve command tool” on page 61.

Installation of major files on your AIX host system

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

Table 8. Major files included in the SDD installation package

Filename	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.

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

addpaths	The command that dynamically adds more paths to SDD devices while they are in `Available` state.
cfgvpath	Configure method of SDD vpath devices.
cfallvpath	Fast-path configure 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 a 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.
mkvg4vp	The command that creates a SDD volume group.
extendvg4vp	The command that extends SDD devices to a SDD volume group.
dpovgfix	The command that fixes a 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.

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 11 lists and describes SDD support for 32-bit and 64-bit applications on AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0.

Table 7 on page 12 lists the various scenarios and the installation packages that you must use for your environment.

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 on 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) or 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 is displayed.
4. Highlight **Install Software** and press Enter.
5. Press F4 to display the INPUT Device/Directory for Software panel.
6. Select the compact disc drive that you are using for the installation, for example, /dev/cd0; and press Enter.
7. Press Enter again. The Install Software panel is displayed.
8. Highlight **Software to Install** and press F4. The Software to Install panel is displayed.
9. Select the installation package that is appropriate for your environment.
10. Press Enter. The Install and Update from LATEST Available Software panel is displayed with the name of the software 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.

Tip:

The published AIX 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, since 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 your currently installed version of SDD

You can verify your previously installed version of SDD by issuing the following command from the root (/) directory:

```
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.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             1.3.2.6  COMMITTED   IBM SDD
                                                  AIX V432 V433 for concurrent
                                                  HACMP

Path: /etc/objrepos
  ibmSdd_432.rte             1.3.2.6  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             1.3.2.6  COMMITTED   IBM SDD
                                                  AIX V433  for nonconcurrent
                                                  HACMP

Path: /etc/objrepos
  ibmSdd_433.rte             1.3.2.6  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             1.3.2.6  COMMITTED   IBM SDD
                                                  AIX V510 for concurrent HACMP

Path: /etc/objrepos
  ibmSdd_510.rte             1.3.2.6  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      1.3.2.6  COMMITTED   IBM SDD
                                                  AIX V510 for nonconcurrent
                                                  HACMP

Path: /etc/objrepos
  ibmSdd_510nchacmp.rte      1.3.2.6  COMMITTED   IBM SDD
                                                  AIX V510 for nonconcurrent
                                                  HACMP
```

Note: The SDD server (sddsrv) is an integrated component of SDD 1.3.2.6 (or later). The SDD server daemon is automatically started after SDD 1.3.2.6 is installed. Go to “Verifying if the SDD server has started” on page 26 to see if the SDD server has automatically started. Go to “SDD server daemon” on page 26 for more details about the SDD server daemon.

Preparing to configure SDD

Before you configure SDD, ensure that:

- The ESS is operational.
- The ibmSdd_nnn.rte software is installed on the AIX host system.
- The ESS hdisks are configured correctly on the AIX host system.

Configure the ESS devices before you configure SDD. If you configure multiple paths to an ESS device, make sure that all paths (hdisks) are in `Available` state. Otherwise, some SDD devices will lose multiple-path capability.

Perform the following steps:

1. Issue the **lsdev -Cc disk | grep 2105** command to check the ESS device configuration.
2. If you have already created some ESS volume groups, vary off (deactivate) all active volume groups with ESS 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 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 (deactivate) all active volume groups with ESS SDD disks.

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 is displayed.
2. Highlight **Data Path Device** and press Enter. The Data Path Device panel is displayed.
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 vpath device configuration” on page 45.
5. Enter the **varyonvg** command to vary on all deactivated ESS volume groups.
6. If you want to convert the ESS hdisk volume group to SDD vpath devices, you must run the hd2vp utility. (See “hd2vp and vp2hd” on page 60 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 (vpath*N*) to ESS devices (hdisks).

Note: If you are running HACMP with 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. See "Special requirements" on page 35.

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 Management (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 -l vpath***N* 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 return it, follow the procedure described in "Configuring SDD" on page 17.

Perform the following steps 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 is displayed.
2. Highlight **Devices** and press Enter.
3. Highlight **Data Path Device** and press Enter. The Data Path Device panel is displayed.
4. Highlight **Remove a Data Path Device** and press Enter. A list of all SDD devices and their state (either `Defined` or `Available`) is displayed.
5. Select the device that you want to unconfigure. Select whether or not you want to delete the device information from the ODM database.
6. Press Enter. The device is unconfigured to the state that you selected.
7. To unconfigure more SDD devices you have to repeat steps 4 - 6 for each SDD device.

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

Verifying the SDD configuration

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

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 is displayed.
2. Select **Data Path Device** and press Enter. The Data Path Device panel is displayed.
3. Select **Display Data Path Device Configuration** and press Enter to display the state (either `Defined` or `Available`) of all SDD pseudo devices and the paths to each device.

If any device is listed as `Defined`, the configuration was not successful. Check the configuration procedure again. See “Configuring SDD” on page 17 for the procedure.

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

1. Type `smitty device` from your desktop window. The Devices menu is displayed.
2. Highlight **Data Path Device** and press Enter. The Data Path Device panel is displayed.
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`)
- The names of the AIX disk devices making up the pseudo device and their configuration and physical volume state

Changing the path-selection policy

SDD 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:

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.

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 path-selection policy is set at the SDD device level. The default path-selection policy for a SDD device is load balancing. You can change the policy for a SDD device with the **chdev** command.

Before changing the path-selection policy, determine the active attributes for the SDD device. Type the `lsattr -El vpathN` command. Press Enter, where *N* represents the vpath number. 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
```

The path-selection policy is the only attribute of an SDD device that can be changed. The valid policies are *rr, lb*, *fo*, and *df*. Here are the explanations for these policy values:

rr round robin

fo failover only

lb load balancing

df (load balancing) default policy

Attention: By changing an SDD device's attribute, the **chdev** command unconfigures and then reconfigures the device. You must ensure that the device is not in use if you are going to change its attribute. Otherwise, the command fails.

Use the following command to change the SDD path-selection policy:

```
chdev -l vpathN -a policy=[rr/fo/lb/df]
```

Adding paths to SDD devices of a volume group

With SDD 1.3.1.3 (or later), you can add more paths to SDD devices that belong to a volume group 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 AIX 4.3.2 or higher, you can issue the **addpaths** command to add paths to SDD devices of a volume group.

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.

You can issue the **addpaths** command to add a new path to a vpath device that has only one existing path. But the new path is not automatically in the `Open` state; 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. See *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 (`vpathN`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. Stop all I/O operations by stopping all applications to the volume group.

   The **addpaths** command is designed to add paths when there are no I/O activities. The command fails if it detects active I/Os.
5. 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 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.
6. Issue the **addpaths** from the AIX command line to add more paths to the SDD devices.
7. 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.

Opening an invalid or close_dead path dynamically

With SDD 1.3.2.6 (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 command” on page 162 in Chapter 8, “Using the datapath commands” on page 153 for more information.

Upgrading SDD for AIX 4.3.2, AIX 4.3.3, and AIX 5.1.0

Understanding SDD support for installation upgrade

Attention: The upgrade installation requires that you:

- Perform a system restart.

With the upgrade installation capability, all configurations are automatically updated during the system restart time. The system restart will start all the automatic configuration, which loads the new driver into the kernel of the host system. The upgrade installation capability is beneficial if you have a large number of SDD devices, volume groups, and file systems configured.

Without the installation upgrade capability, you must *manually* upgrade to a *newer* version of SDD. To manually upgrade to a newer version of SDD, see “Upgrading SDD manually” on page 23.

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

SDD 1.3.0.1 (or later) supports installation upgrade if you are upgrading from any one of the installation packages listed in Table 9.

Table 9. List of previously installed installation packages that are supported with the installation upgrade

Installation package name
ibmSdd_432.rte
ibmSdd.rte.432
ibmSdd_433.rte
ibmSdd.rte.433
ibmSdd_510.rte
ibmSdd_510nchacmp.rte

If you previously installed from any of the installation packages listed in Table 9, with SDD 1.3.0.1 (or later) you can upgrade while all of the SDD:

- File systems are mounted
- Volume groups are varied-on

If you are upgrading from a previous version of SDD that you installed from other installation packages, you cannot perform the installation upgrade. To upgrade SDD to a newer version, all the SDD installation packages must be uninstalled.

To verify your currently installed version of SDD , go to “Verifying your currently installed version of SDD” on page 15.

If the previous version of SDD is installed from one of the installation packages listed in Table 9, proceed to “Upgrading SDD” on page 23.

If the previous version of SDD is installed from a installation package *not* listed in Table 9, proceed to “Removing SDD from an AIX host system” on page 24 and “Installing SDD” on page 14.

Verifying your previously installed version of SDD

You can verify your previously installed version of SDD by issuing the following command from the root (/) directory:

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

(where 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 the following commands:

```
lslpp -l ibmSdd_432.rte

lslpp -l ibmSdd.rte.432

...
```

...

Upgrading SDD

Note: The SDD server (sddsrv) is an integrated component of SDD 1.3.2.6 (or later). The SDD server daemon is automatically started after SDD is installed. You must stop the SDD server if it is running in the background *before proceeding with the upgrade instructions.* Go to "Verifying if the SDD server has started" on page 26 and "Stopping the SDD server" on page 27 for more instructions. Go to "SDD server daemon" on page 26 for more details about the SDD server daemon.

You can use SDD 1.3.0.1 (or later) using an installation upgrade if you are upgrading from any of the installation packages listed in Table 9 on page 22. See "Understanding SDD support for installation upgrade" on page 21 for more information.

Perform the following steps to upgrade SDD with an installation upgrade:

1. Complete the installation instructions provided in "Installing SDD" on page 14.
2. Restart your system by typing `shutdown -rF`
3. Verify the SDD configuration by typing `lsvpcfg`
4. Verify your currently installed version of the SDD by completing the instructions provided in "Verifying your currently installed version of SDD" on page 15

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

Upgrading SDD manually

If you are upgrading from a previous version of the SDD that you installed with a installation package not listed in Table 9 on page 22, you cannot perform an installation upgrade.

Perform the following steps to upgrade SDD:

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

   ```
   rm .toc
   ```

   Note: Ensure that this file is removed because it contains information about the previous version of SDD or DPO.
2. Issue the **lspv** command to find out all the SDD volume groups.
3. Issue the **lsvgfs** command for each SDD volume group to find out which file systems are mounted. Type the following command:

   ```
   lsvgfs  vg_name
   ```
4. Issue the **umount** command to unmount all file systems belonging to SDD volume groups. Type the following command:

   ```
   umount  filesystem_name
   ```
5. Run the **vp2hd** script to convert the volume group from SDD devices to ESS hdisk devices.
6. Issue the **varyoffvg** command to vary off the volume groups. Type the following command:

   ```
   varyoffvg  vg_name
   ```

7. Remove all SDD devices. Type the following command:

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

8. Use the **smitty** command to uninstall SDD. Type `smitty deinstall` and press Enter. The uninstall process begins. Complete the uninstall process. See "Removing SDD from an AIX host system" for the step-by-step procedure for uninstalling SDD.
9. Use the **smitty** command to install the newer version of SDD from the compact disc. Type `smitty install` and press Enter. The installation process begins. Go to "Installing SDD" on page 14 to complete the installation process.
10. Use the **smitty device** command to configure all the SDD devices to the `Available` state. See "Configuring SDD" on page 17 for a step-by-step procedure for configuring devices.
11. Issue the **lsvpcfg** command to verify the SDD configuration. Type the following command:

    ```
    lsvpcfg
    ```

12. Issue the **varyonvg** command for each volume group that was previously varied offline. Type the following command:

    ```
    varyonvg vg_name
    ```

13. Run the **hd2vp** script for each SDD volume group to convert the physical volumes from ESS hdisk devices back to SDD vpath devices. Type the following command:

    ```
    hd2vp  vg_name
    ```

14. Issue the **lspv** command to verify that all physical volumes of the SDD volume groups are SDD vpath devices.

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

Removing SDD from an AIX host system

Note:

- The SDD server (sddsrv) is an integrated component of SDD 1.3.2.6 (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 26 and "Stopping the SDD server" on page 27 for more instructions. Go to "SDD server daemon" on page 26 for more details about the SDD server daemon.
- Before you remove the SDD package from your AIX host system, all the SDD devices must be unconfigured and removed from your host system.
- See "Unconfiguring Subsystem Device Driver" on page 17.

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:

   ```
   ibmSdd_432.rte
   ```

```
ibmSdd_433.rte
ibmSdd_510.rte
ibmSdd_510nchacmp.rte
```

Then press Enter.

Note: See "Verifying your currently installed version of SDD" on page 15 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 or version of SDD.

You could also press F4 in the SOFTWARE name field 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.

Using ESS concurrent download of licensed internal code

Concurrent download of licensed internal code 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: During the download of licensed internal code, the AIX error log might overflow and excessive system paging space could be consumed. When the system paging space drops too low it could cause your AIX system to hang. To avoid this problem, you can perform the following steps prior to doing the download:

1. Save the existing error report by typing the following command from the AIX command-line interface:

   ```
   > errpt > file.save
   ```
2. Delete the error log from the error log buffer by typing the following command:

   ```
   > errclear 0
   ```
3. Enlarge the system paging space by using the SMIT tool.
4. Stop the AIX error log daemon by typing the following command:

   ```
   /usr/lib/errstop
   ```

Once you have completed steps 1- 4, you can perform the download of the ESS licensed internal code. After the download completes, type `/usr/lib/errdemon` from the command-line interface to restart the AIX error log daemon.

SDD server daemon

The SDD server (also referred to as sddsrv) is an integrated component of SDD 1.3.2.6 (or later). This component consists of a UNIX application daemon that is installed in addition to the SDD device driver. The SDD server daemon starts automatically after the SDD driver package is installed. The SDD server daemon runs in the background at all times. The daemon scans to find failing paths (INVALID, CLOSE_DEAD, or DEAD) regularly between 2-5 minutes intervals. The daemon scans CLOSE paths and OPEN paths which are idle (failed paths) more often at 1 minute regular intervals.

Understanding how the SDD server daemon works

The SDD server 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 - 5 minute intervals between consecutive runs. If the test succeeds then SDD server reclaims these paths , and changes the states of these paths according to the following characteristics:

- If the state of the vpath is OPEN then the SDD server changes the state of INVALID and CLOSE_DEAD paths of that vpath to OPEN.
- If the state of the vpath is CLOSE then the SDD server changes the state of INVALID and CLOSE_DEAD paths of that vpath to CLOSE.
- The SDD server changes the state of DEAD paths to OPEN.

Path probing

The SDD server will regularly test CLOSE paths and OPEN paths which are idle to see if they are working or has become inoperational. The daemon sleeps for 1 minute intervals between consecutive runs. It will test OPEN and CLOSE paths, and detect if these paths have become idle. If the test fails, the SDD server then changes the state of these paths according to the following characteristics:

- If the vpath is in the OPEN state and the path is not working, then SDD server changes the state of the path from OPEN to DEAD.
- If vpath is in the CLOSE state and the path is not working, then SDD server changes the state of the path from CLOSE to CLOSE_DEAD.

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 27 to proceed.

Starting the SDD server manually

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

Go to "Verifying if the SDD server has started" on page 26 to see If you successfully started the SDD server.

Note: The **startsrc –s 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" 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 p`*PN* `–s sddsrv` 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 sddsrv`
2. Type `ps -ef | grep sddsrv` 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 26 to determine if the SDD servers state is *Active*.

   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

You can stop the SDD server by typing `stopsrc -s sddsrv`. After issuing this command, you need to verify if the SDD server has stopped. See "Verifying if the SDD server has started" on page 26 and the status should be *inoperative*.

Understanding SDD 1.3.2.6 (or later) support for single-path configuration

SDD 1.3.2.6 (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: SDD supports one SCSI adapter on the host system.

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

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

Tip: It is also possible to change from single-path to multi-path configuration by using the **addpaths** command. For more information about the **addpaths** command, go to "Adding paths to SDD devices of a volume group" on page 20.

Understanding SDD error recovery policies

There are differences in the way that SDD 1.3.2.6 (or later) and SDD 1.3.1.3 (or earlier) handles error recovery for AIX host systems:

SDD 1.3.1.3 (or earlier) error recovery policy
With SDD 1.3.1.3 (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 to become disabled in an event of transient errors occurring. 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 1.3.2.6 (or later) error recovery policy
With SDD 1.3.2.6 (or later), error recovery policy is designed to report errors to applications more quickly. With SDD 1.3.2.6 (or later), the applications receive failed I/O requests quicker. 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 1.3.2.6 (or later) and SDD 1.3.1.3 (or earlier) error recovery policies support the following modes of operation:

single-path mode
An AIX 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 2 times.
- With the SDD 1.3.2.6 (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 1.3.2.6 (or later) the SDD server daemon detects the single CLOSE path that is failing it will change 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.

multi-path mode
The host system has multiple paths that are configured to an ESS LUN.

Both the SDD 1.3.2.6 (or later) and SDD 1.3.1.3 (or earlier) error recovery policies in multiple-path mode have the following common 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 on an operational path. SDD 1.3.2.6 (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.
- 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 1.3.2.6 (or later) error recovery policy in *multiple-path mode* has the following latest characteristics:

- 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 command” on page 161 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 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 *multiple-path mode* has the following characteristics:

- 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 the SDD support for pSeries 690 with static LPARs configured

pSeries 690 supports static LPARs as a standard feature, and any customer 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 his its own instance of a operating system and does not share pSeries hardware resources with any other partition. So each partition functions the same as it does on a standalone system. Storage subsystems need to be shared the same way 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, then the path optimization can be performed on those adapters as the same way in a standalone system. When the adapters are not shared with any other partition, SCSI reservation, persistent reserve, and LUN level Masking operates as expected (by being ″bound″ to an instance of the operating system). It is recommended to install AIX SDD persistent reserve version for AIX 5.1.0 (ibmSdd_510nchacmp.rte), when multiple LPARs are configured or are planned to be configured in the future. If the plan is to share storage subsystem devices in a nonconcurrent mode between multiple LPARs on a pSeries 690 system, then the user must install the ibmSdd_510nchacmp.rte package. Failure to use this persistent reserve version of SDD with this configuration will expose the disks to possible data corruption. This statement is true with and without HACMP installed on a pSeries 690 system.

Understanding the SDD support for High Availability Cluster Multi-Processing (HACMP)

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. RS/6000 / pSeries servers running HACMP in concurrent or nonconcurrent mode are supported. Different SDD releases support different kinds of environments. (See Table 10 on page 31, Table 12 on page 32, Table 11 on page 32 and Table 13 on page 33.)

HACMP provides a reliable way for clustered IBM RS/6000 / pSeries servers which share disk resources to recover from server and disk failures. In a HACMP environment, each RS/6000 / pSeries server in a cluster is a node. Each node has access to shared disk resources that are accessed by other nodes. 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.

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, in order 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 SDD 1.3.2.6 (or later), or the ibmSdd_510.rte installation package for SDD 1.3.2.6 when running HACMP concurrent mode.

- Use either ibmSdd_433.rte installation package for SDD 1.3.2.6 (or later), or the ibmSdd_510nchacmp.rte installation package for SDD 1.3.2.6 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:

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

SDD supports RS/6000 / 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. Table 10 and Table 12 on page 32 summarizes the software requirements to support HACMP:

Note: The following tables still needs to be updated.

Table 10. 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)	• IY07392 • IY03438 • IY11560 • IY08933 • IY11564 • IY12021 • IY12056 • F model requires IY11110	• IY11563 • IY11565 • IY12022 • IY12057 • F model requires IY11480	Not applicable.
ibmSdd_432.rte installation package for SDD 1.2.0.0 and 1.2.2.X (SCSI/FCP)	• IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • F model requires IY11110	• IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • F model requires IY11480	Not applicable.
ibmSdd_432.rte installation package for SDD 1.3.1.3, SDD 1.3.2.6 (SCSI/FCP)	• IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • F model requires IY11110	• IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • F model requires IY11480	• (See Note.)
Note: For up-to-date APAR information for HACMP, see the following Web site: http://hacmp.aix.dfw.ibm.com			

Table 11. 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)	• IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • IY14682 • F model requires IY11110	• IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • IY14683 • F model requires IY11480	Not applicable.
ibmSdd_433.rte installation package for SDD 1.3.1.3 and SDD 1.3.2.6 (SCSI/FCP)	• IY07392 • IY13474 • IY03438 • IY08933 • IY11560 • IY11564 • IY12021 • IY12056 • IY14682 • F model requires IY11110	• IY13432 • IY11563 • IY11565 • IY12022 • IY12057 • IY14683 • F model requires IY11480	• (See Note.)
Note: For up-to-date APAR information for HACMP, see the following Web site: http://hacmp.aix.dfw.ibm.com			

Table 12. 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)
ibmSdd_510.rte installation package for SDD 1.3.1.3 (SCSI/FCP)	• IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480	• IY23761 • IY26561 • (See Note.)

Table 12. 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)
ibmSdd_510.rte installation package for SDD 1.3.2.6 (SCSI/FCP_	• IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480	• IY23761 • IY26561 • IY25995 • (See Note.)
Note: For up-to-date APAR information for HACMP, see the following Web site: http://hacmp.aix.dfw.ibm.com		

Table 13. 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)
ibmSdd_510nchacmp.rte installation package for SDD 1.3.1.3 (SCSI/FCP)	• IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480	• IY23761 • IY26561 • (See Note.)
ibmSdd_510nchacmp.rte installation package for SDD 1.3.2.6 (SCSI/FCP)	• IY11563 • IY11565 • IY12022 • IY12057 • IY13432 • IY14683 • IY17684 • IY19089 • IY19156 • F model requires IY11480	• IY23761 • IY26561 • IY25995 • (See Note.)
Note: Only HACMP 4.4.1 and APARs support 64–bit mode. For up-to-date APAR information for HACMP, see the following Web site: http://hacmp.aix.dfw.ibm.com		

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

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

- service.software.ibm.com/rs6k/fixdb.html

Even though SDD supports HACMP, certain combinations of features are not supported. Table 14 lists those combinations:

Table 14. HACMP and supported SDD features

Feature	RS/6000 / pSeries node running HACMP
ESS concurrent code load on multi-path 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

SDD installation package attributes

The ibmSdd_433.rte and ibmSdd_510nchacmp.rte installation packages for SDD 1.3.2.0 (or later) have different features compared with ibmSdd_432.rte and ibmSdd_510.rte installation packages for SDD 1.3.2.0 (or later). The ibmSdd_433.rte and ibmSdd_510nchacmp.rte installation packages implement the SCSI-3 Persistent Reserve command set, in order to support HACMP in nonconcurrent mode with single-point failure protection. The ibmSdd_433.rte and ibmSdd_510nchacmp.rte installation packages 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 ibmSdd_433.rte and ibmSdd_510nchacmp.rte installation packages will switch the multipath configuration to a single-path configuration. There is no single-point failure protection for single-path configurations.

The ibmSdd_433.rte and ibmSdd_510nchacmp.rte installation packages have a new attribute under its pseudo parent (dpo), that reflects whether the ESS supports the Persistent Reserve Command set or not. 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 its value is set to `yes`; otherwise this attribute only exists in the PdAt ODM and its 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 attached ESS has 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. There are 2 ways to get a unique reservation key. 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 generates a unique reservation key by using the 6 digits of the machine ID from issuing the **uname -m** command and appending 2 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 configure methods generates a unique reservation key by appending 2

digits of LPAR ID from the output of the **uname -L** command to the 6 digits of the machine ID from issuing the **uname -m** command.

To check the Persistent Reserve Key of a node, provided by HACMP, 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 the ibmSdd_433.rte and ibmSdd_510nchacmp.rte installation packages for the SDD 1.3.2.0 (or later) vpath devices. You must unconfigure and remove the vpath devices *before* you unconfigure and remove the vpath devices' underlying ESS hdisks. Otherwise, if the ESS 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; however, SDD does not. 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 59) to convert the volume group's physical volumes from ESS hdisks to vpath devices. This conversion step not only creates *pvid* attributes for all vpath devices which 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.

Under certain conditions, the state of a physical device's *pvid* on a system is not always as expected. So 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 pvid's for both hdisks and vpath:

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

Scenario 2. lspv displays pvid's 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 48 for a detailed procedure for importing a volume group with the SDD devices.

Table 15. HACMP and supported features

Feature	**RS/6000 / pSeries node in HACMP nonconcurrent mode**
ESS concurrent code load	Yes
SDD load balancing	Yes
SCSI	Yes
FCP (fibre)	Yes
Single-path fibre	No
SCSI and single-path fibre (mixed environment)	No

Table 16. HACMP and supported features

Feature	**RS/6000 / pSeries node in HACMP nonconcurrent mode**
ESS concurrent code load	Yes
SDD load balancing	Yes
SCSI	Yes
FCP (fibre)	Yes
Single-path fibre	No
SCSI and single-path fibre (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 http://www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates

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

With AIX v5.1.D and HACMP v4.4.1.4, you can create enhanced concurrent capable volume groups with ESS 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.

Note: There are different ways for creating and varying-on concurrent capable volume groups, however to extend, reduce, and export a volume group, the instructions are the same.

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 provides information and instructions on the operating actions the user can perform.

Creating HACMP RAID concurrent mode volume groups

Perform the following steps to create a HACMP RAID concurrent mode volume group:

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

1. Type smitty `datapath_mkvg` at the command prompt.

A screen similar to the following is displayed. The following example is 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 desired 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]
*******************************************************************************
```

Importing HACMP RAID concurrent mode volume groups

When importing the volume group to other nodes in the cluster, the customer 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 a number of ways:

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

- 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>
  ```

  And 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)
  ```

  Note: If the hdisks do not have pvid, issuing the **chdev -l hdiskX -a pv=yes** command should recover it. Verify pvid exists by issuing the **lspv** command against the hdisk. Same method can be used when attempting to obtain pvid on vpath.

  Note: You should also verify that the volume group is not varied-on on 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>
```

- 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)
  zabra>

Note: There is a possibility that the vpath value may be different for each host. You will need to verify that the value (21B21411) following the vpath name are identical on each node.

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

1. Type `smitty importvg` at the command prompt.

Note: A screen similar to the following is displayed. The following example displays how to import a 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?          yes
Make default varyon of VOLUME GROUP concurrent?     no
********************************************************************************
```

Once importing of volume groups has been completed, issue the **lsvpcfg** command again to verify the state of the vapth.

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

Then, 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, *its 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, it requires a special option flag *-u* 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 a HACMP RAID concurrent mode volume group:

Note: Removing a HACMP RAID concurrent mode volume group can be accomplished by exporting volume groups, or by following the procedure below.

1. Type `smitty vg` at the command prompt.
2. Select *Remove a Volume Group* from the displayed menu.

Note: A screen similar to the following is displayed. The following example displays how to remove a 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 a HACMP RAID concurrent mode volume groups.

1. Type `smitty datapath_extendvg` at the command prompt.

Note: A screen similar to the following is displayed. The following example displays how to extend a 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]
********************************************************************************
```

Note: Verify that extending 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.

Reducing HACMP RAID concurrent mode volume groups

Perform the following steps to reduce a HACMP RAID concurrent mode volume group.

1. Type `smitty vg` at the command prompt.
2. Select *Set Characteristics of a Volume Group* from the displayed menu.
3. Select *Remove a Physical Volume from a Volume Group* from the displayed menu.

Note: A screen similar to the following is displayed. The following example displays how to reduce a HACMP RAID concurrent mode 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.

Exporting HACMP RAID concurrent mode volume groups

Perform the following steps to export a HACMP RAID concurrent mode volume group.

1. Type `exportvg` at the command prompt.

   Note: To export *con_vg*, use the **exportvg con_vg** command.

   Note: Before exporting a HACMP RAID concurrent mode volume group, make sure the volume group is varied-off.

Enhanced concurrent capable volume groups

With SDD 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 the user creates 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), is propagated to all other nodes. Also, the logical volume manager (LVM) configuration files are updated on all nodes.

The following sections provides information and instructions on the operating actions the user can perform.

For more detailed information on enhanced concurrent capable volume groups, see “Supporting enhanced concurrent mode in an HACMP environment” on page 43.

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 what the next available major number is. The volume groups must be created with a major number that is available on all nodes. See the following for an example:

```
dollar>lvlstmajor
41,54..58,60..62,67,78...
dollar>
```

```
monkey>lvlstmajor
39..49,55,58,67,80...
monkey>
zebra>lvlstmajor
40..49,55,58..61,67,78...
zebra>
```

From the above 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 customer may choose to begin with the highest available (80) and increment from there.

1. Type `smitty datapath_mkvg` at the command prompt.

   Note: A screen similar to the following is displayed. The following example displays 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                          [45]
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 steps to import enhanced concurrent capable volume groups.

1. Type `smitty importvg` at the command prompt.

Note: Before importing enhanced concurrent capable volume groups on SDD devices, make sure there is pvid on the vpath. This can be displayed by typing the **lspv** command. If pvid is not displayed, you need to import the volume group on one of the hdisks that belongs to the vpath. Type `hd2vp` to convert the volume group to vpath based.

Note: If the hdisks do not have a pvid, issuing the **chdev -l hdisk -a pv=yes** should recover it. Verify that pvid now exists by issuing the **lspv** command against the hdisk. This same method can be used when attempting to obtain a pvid on a vpath.

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

Note: A screen similar to the following is displayed. The following example displays 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?          no
Make default varyon of VOLUME GROUP concurrent?        no
Importing a concurrent volume group.          no
********************************************************************************
```

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

For more information, detailed procedures and examples of enhanced concurrent capable volume groups, see the following section.

- "Removing HACMP RAID concurrent mode volume groups" on page 40
- "Extending HACMP RAID concurrent mode volume groups" on page 40
- "Exporting HACMP RAID concurrent mode volume groups" on page 41

Recovering paths that are lost during HACMP node failover

Normally, when there is a node failure, HACMP transfers ownership of shared disks and other resources, through a process known as node failover. Certain situations, such as a loose or disconnected SCSI or fibre-adapter card, can cause your vpath devices to lose one or more underlying paths during node failover. 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 "Adding paths to SDD devices of a volume group" on page 20 for more information about the **addpaths** command.

Note: Simply running the **cfgmgr** command while the vpath devices are in the `Available` state will not recover the lost paths; that is why you need to 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. You also need bos.clvm.enh and bos.rte.lvm filesets installed at level 5.1.0.10 on all the nodes.

SDD 1.3.2.6 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 -c (in 32-bit kernel) or -C (in 64-bit kernel) option for **mkvg4vp** command. To create enhanced concurrent capable volume groups from smit panel, the user needs to set `yes` for **Create Volume Group Concurrent Capable?**. Both ways will

leave the enhanced concurrent capable volume group in varied-off mode. The user should import this concurrent volume group to all other nodes and add the volume group into the HACMP concurrent resource group, then start the HACMP. The volume group will be varied-on by HACMP. After the changes are made to one volume group on one node, it will automatically propagate all changes 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 41.

Managing secondary system paging space

SDD 1.3.2.6 (or later) supports secondary system paging on *multi-path* fibre-channel vpath device from an AIX 4.3.3 or AIX 5.1.0 host system to an ESS.

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

Note: SDD supports secondary paging on ESS. The benefits are multi-pathing to your paging spaces. All the same commands apply to using vpath-based volume groups for paging spaces as hdisk-based volume groups.

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 start-up, which can result in poor start-up performance. **Do not** redefine your primary paging space using vpath devices.

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

http://techlib.austin.ibm.com:80/techlib/manuals/adoclib/aixbman/baseadmn/pagoverv.htm

Listing paging spaces

You can list paging spaces by typing the following:

```
lsps -a
```

Adding a paging space

You can add a paging space by typing the following:

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

The **mkps** 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*NN* The –s argument accepts the number of logical partitions (*NN*) to allocate to the new paging space.

vg *vg* 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.

For example, to remove paging space PS01, type:

```
rmps PS01
```

Providing load-balancing and failover protection

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

Displaying the ESS vpath device configuration

To provide failover protection, an ESS vpath device must include a minimum of two paths. Both the SDD vpath device and the ESS 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 ESS LUN. The other SDD vpath devices have two paths and, therefore, can provide failover protection.

To display which ESS 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 is displayed.
2. Select **Data Path Devices** and press Enter. The Data Path Devices panel is displayed.
3. Select **Display Data Path Device Configuration** and press Enter.

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 )
```

Figure 2. Output from the Display Data Path Device Configuration SMIT panel

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 ESS vpath devices each have two paths and thus can provide failover protection. The requirement for failover protection is that the ESS 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 ESS hdisk devices, 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 LUN, such as `019FA067`
- The names of the AIX disk devices that comprise the SDD vpath devices, their configuration state, and the physical volume state.

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   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.

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 Figure 2 on page 45) and then add paths later by reconfiguring the ESS. (See “Adding paths to SDD devices of a volume group” on page 20 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 only has a single path.

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

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

   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 **System Storage Management (Physical & Logical Storage)** and press Enter. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press Enter. The Volume Groups panel is displayed.

4. Select **Volume Group** and press Enter. The Add Volume Groups with Data Path Devices panel is displayed.
5. Select **Add a Volume Group with Data Path Devices** and press Enter.

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

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

All the functions that apply to a regular volume group also apply to an SDD volume group. Use SMIT to create a logical volume group (mirrored, striped, or compressed) or a file system (mirrored, striped, or compressed) on an SDD volume group.

Once you create the volume group, AIX creates the SDD vpath device as a physical volume (pv). In Figure 2 on page 45, 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 ESS 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 is displayed.

   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 **System Storage Management (Physical & Logical Storage)** and press Enter. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press Enter. The Volume Group panel is displayed.
4. Select **Volume Groups** and press Enter. The Volume Groups panel is displayed.
5. Select **Import a Volume Group** and press Enter. The Import a Volume Group panel is displayed.
6. In the Import a Volume Group panel, perform the following tasks:
   - Type in the volume group you want to import.
   - Type in the physical volumes that you want to import over.
   - Press Enter after making all desired 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 59) to convert the volume group's physical volumes from ESS hdisks to vpath devices. This conversion step not only creates *pvid* attributes for all vpath devices which 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 35 for special requirements regarding unconfiguring and removing the ibmSdd_433.rte or ibmSdd_510nchacmp.rte installation packages for SDD 1.3.2.0 (or later) vpath devices.

See 36 for a scenario case on when hdisks' PVID do not exist.

Exporting a volume group with SDD

You can export a volume group definition from a set of physical volumes 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. Since 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 is displayed.

   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 **System Storage Management (Physical & Logical Storage)** and press Enter. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press Enter. The Volume Group panel is displayed.
4. Select **Volume Groups** and press Enter. The Volume Groups panel is displayed.
5. Select **Export a Volume Group** and press Enter. The Export a Volume Group panel is displayed.
6. Type in the volume group to export and press Enter.

You can use the F4 key to select which volume group you want to export.

Losing failover protection

AIX can only create volume groups 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 only has 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 devices does not have failover protection because there is no alternate path to the ESS 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 desired pvid to activate the volume group.

As an example, if you issue the **lsvpcfg** command, the following output is displayed:

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

The following example of a **chdev** command could also set the pvid attribute for an hdisk:

```
chdev -l hdisk46 -a queue_depth=30
```

For this example, the output of the **lsvpcfg** command would look similar to this:

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

The output of the **lsvpcfg** command shows that vpath11 contains hdisk22 and hdisk46. However, hdisk46 is the one with the pv attribute set. If you run the **lsvg -p vpathvg** command again, you might see something like this:

```
vpathvg:
PV_NAME           PV STATE    TOTAL PPs   FREE PPs    FREE DISTRIBUTION
vpath10           active      29          4           00..00..00..00..04
hdisk46           active      29          4           00..00..00..00..04
vpath12           active      29          4           00..00..00..00..04
vpath13           active      29          28          06..05..05..06..06
```

Notice that now device vpath11 has been replaced by hdisk46. That is because hdisk46 is one of the hdisk devices included in vpath11 and it has a pvid attribute in the ODM. In this example, the LVM used hdisk46 instead of vpath11 when it activated volume group vpathvg. The volume group is now in a mixed mode of operation because it partially uses vpath pseudo devices and partially uses hdisk devices. This is a problem that must be fixed because failover protection is effectively disabled for the vpath11 physical volume of the vpathvg volume group.

Note: The way to fix this problem with the mixed volume group is to run the **dpovgfix vg-name** command after running the **chdev** command.

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

Assume that vpath3 is made up of hdisk4 and hdisk27 and that vpath3 is currently a physical volume. If the vpath3, hdisk4, and hdisk27 devices are all deleted by using the **rmdev** command and then **cfgmgr** is invoked at the command line, only one path of the original vpath3 is configured by AIX. The following commands would produce this situation:

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

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

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

- 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 devices is not always required. It depends on whether the ESS 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. Since the ESS 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 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 device has never belonged to a volume group, that means there is no pvid written on its sector 0. In that case, you only need to run **cfgmgr** command 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.

  Note: 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 “Adding paths to SDD devices of a volume group” on page 20 section 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 tries to find 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 Figure 2 on page 45) and then add paths later by reconfiguring the ESS. With a single path, failover protection is not provided. (See “Adding paths to SDD devices of a volume group” on page 20 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 is displayed.

   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 **System Storage Management (Physical & Logical Storage)** and press Enter. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press Enter. The Volume Group panel is displayed.
4. Select **Volume Group** and press Enter. The Add Volume Group with Data Path Devices panel is displayed.

5. Select **Add a Data Path Volume to a Volume Group** and press Enter.
6. Type in 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 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 a 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 backup a volume group with SDD devices, go to “Accessing the Back Up a Volume Group with Data Path Devices SMIT panel” on page 58.

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 backup. Keep system activity at a minimum during the system backup procedure.

Restoring all files belonging to a 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 and go to “Accessing the Remake a Volume Group with Data Path Devices SMIT panel” on page 59.

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 requires 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 17 lists the SDD-specific SMIT panels and how you can use them.

Table 17. SDD-specific SMIT panels and how to proceed

SMIT panels	How to proceed:
Display Data Path Device Configuration	Go to: “Accessing the Display Data Path Device Configuration SMIT panel” on page 55
Display Data Path Device Adapter Status	“Accessing the Display Data Path Device Adapter Status SMIT panel” on page 55
Display Data Path Device Adapter Status	“Accessing the Display Data Path Device Adapter Status SMIT panel” on page 55
Define and Configure all Data Path Devices	“Accessing the Define and Configure All Data Path Devices SMIT panel” on page 56

Table 17. SDD-specific SMIT panels and how to proceed (continued)

Add Paths to Available Data Path Devices	"Accessing the Add Paths to Available Data Path Devices SMIT panel" on page 56
Configure a Defined Data Path Device	"Accessing the Configure a Defined Data Path Device SMIT panel" on page 56
Remove a Data Path Device	"Accessing the Remove a Data Path Device SMIT panel" on page 57
Add a Volume Group with Data Path Devices	"Accessing the Add a Volume Group with Data Path Devices SMIT panel" on page 57
Add a Data Path Volume to a Volume Group	"Accessing the Add a Data Path Volume to a Volume Group SMIT panel" on page 57
Remove a Physical Volume from a Volume group	"Accessing the Remove a Physical Volume from a Volume group SMIT panel" on page 58
Back Up a Volume Group with Data Path Devices	"Accessing the Back Up a Volume Group with Data Path Devices SMIT panel" on page 58
Remake a Volume Group with Data Path Devices	"Accessing the Remake a Volume Group with Data Path Devices SMIT panel" on page 59

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

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

   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 is displayed.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Configuration** and press Enter.

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 `smitty` from your desktop window. The System Management Interface Tool is displayed.

   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 is displayed.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Adapter Status** and press Enter.

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

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

   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 is displayed.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Adapter Status** and press Enter.

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

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

   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 is displayed.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel is displayed.
4. Select **Define and Configure All Data Path Devices** and press Enter.

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

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

   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 is displayed.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel is displayed.
4. Select **Add Paths to Available Data Path Devices** and press Enter.

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

1. Type `SMITTY` from your desktop window. The System Management Interface Tool is displayed.

   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 is displayed.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel is displayed.
4. Select **Configure a Defined Data Path Device** and press Enter.

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

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

   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 is displayed.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel is displayed.
4. Select **Remove a Data Path Device** and press Enter.

Accessing the Add a Volume Group with Data Path Devices SMIT panel: Perform the following steps to access the Add a volume group with data path devices panel:

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

   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 **System Storage Management (Physical & Logical Storage)** and press Enter. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press Enter. The Volume Group panel is displayed.
4. Select **Volume Groups** and press Enter. The Add Volume Group with Data Path Devices panel is displayed.
5. Select **Add Volume Group with Data Path Devices** and press Enter.

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

Accessing the Add a Data Path Volume to a Volume Group SMIT panel: Perform the following steps to access the Add a Data Path Volume to a Volume Group panel:

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

   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 **System Storage Management (Physical & Logical)** and press Enter. The System Storage Management (Physical & Logical) panel is displayed.
3. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel is displayed.
4. Select **Volume Group** and press Enter. The Volume Group panel is displayed.
5. Select **Add a Data Path Volume to a Volume Group** and press Enter.
6. 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 is displayed.

   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 **Logical Volume manager** and press Enter. The Logical Volume manager panel is displayed.
3. Select **Volume Groups** and press Enter. The Volume Groups panel is displayed.
4. Select **Set Characteristics of a Volume Group** and press Enter. The Set Characteristics of a Volume Group panel is displayed.
5. Select **Remove a Physical Volume from a Volume group** and press Enter. The Remove a Physical Volume from a Volume Group panel is displayed.

Accessing the 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 backup a volume group with SDD devices:

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

   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 **System Storage Management (Physical & Logical Storage)** and press Enter. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press Enter. The Volume Group panel is displayed.
4. Select **Volume Groups** and press Enter. The Volume Groups panel is displayed.
5. Select **Back Up a Volume Group with Data Path Devices** and press Enter. The Back Up a Volume Group with Data Path Devices panel is displayed.
6. In the Back Up a Volume Group with Data Path Devices panel, perform the following steps:
   - Type in the Backup DEVICE or FILE name.
   - Type in the Volume Group to back up.
   - Press Enter after making all desired changes.

   Tip: 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 backup.

   Attention: Backing-up files (running the **savevg4vp** command) will result in the loss of all material previously stored on the selected output medium. Data integrity of the archive may be compromised if a file is modified during system backup. Keep system activity at a minimum during the system backup procedure.

Accessing the Remake a Volume Group with Data Path Devices SMIT panel:
Perform the following steps to access the Remake a Volume Group with Data Path Devices panel and restore a volume group with SDD devices:

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

   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 **System Storage Management (Physical & Logical Storage)** and press Enter. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press Enter. The Volume Group panel is displayed.
4. Select **Volume Groups** and press Enter. The Volume Groups panel is displayed.
5. Select **Remake a Volume Group with Data Path Devices** and press Enter. The Remake a Volume Group with Data Path Devices panel is displayed.
6. Type in the Restore DEVICE or FILE name, 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.

For more information about this command, go to "Adding paths to SDD devices of a volume group" on page 20.

hd2vp and vp2hd: SDD provides two conversion scripts, hd2vp and vp2hd. The hd2vp script converts a volume group from ESS hdisks into SDD vpaths, and the vp2hd script converts a volume group from SDD vpaths into ESS hdisks. Use the vp2hd program when you want to configure your applications back to original ESS hdisks, or when you want to remove the SDD from your AIX host system.

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

The syntax for these conversion scripts is as follows:

```
hd2vp  vgname
vp2hd  vgname
```

These two conversion programs require that a volume group contain either *all* original ESS 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 hdisk devices can cause original ESS hdisks to be contained within a SDD volume group. This is known as a mixed volume group. Mixed volume groups happen when a 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 hdisk devices. The result is the volume group loses path failover capability for that physical volume. The dpovgfix script tool fixes this problem. The command syntax is:

```
dpovgfix vg-name
```

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 two ways.

1. The command can be issued without parameters. The command syntax is:

```
lsvpcfg
```

See “Verifying the SDD configuration” on page 18 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 18 for an explanation of the output.

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 46.

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 53.

Persistent reserve command tool

With SDD 1.3.2.6 (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.

This command can be issued to both SDD vpath devices and 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 primary resource owner suddenly goes down without releasing the persistent reserve, and for some reason it cannot be brought up for awhile then standby node of LPAR or server cannot take the ownership of sharing resources.

Note: 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 also it 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 Remove 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 devices directly

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

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

Two types of applications use ESS disk storage. One type of application accesses ESS devices through the SDD vpath device (raw device). The other type of application accesses ESS 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 hdisk device special files directly before installing SDD, convert it to using the SDD vpath device special files. After installing SDD, perform the following steps:

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

   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 is displayed.
3. Select **Data Path Devices** and press Enter. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Configuration**. The system displays all SDD vpaths with their attached multiple paths (hdisks).
5. Search the list of hdisks to locate the hdisks your application is using.
6. 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 devices through AIX LVM

Attention:

- You must use the System Management Interface Tool (SMIT). The SMIT facility runs in two interfaces, nongraphical (type `smitty` 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 devices through LVM, determine the volume group that it uses before you convert volume groups. Then, perform the following steps to convert the volume group from the original ESS device hdisks to the SDD vpaths:

1. Determine the file systems or logical volumes that your application accesses.
2. Type `smitty` from your desktop window. The System Management Interface Tool is displayed.
3. Select **System Storage Management (Physical & Logical Storage)** and press Enter. The System Storage Management (Physical & Logical Storage) panel is displayed.
4. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel is displayed.
5. Select **Logical Volume** and press Enter. The Logical Volume panel is displayed.
6. Select **List All Logical Volumes by Volume Group** to determine the logical volumes that belong to this volume group and their logical volume mount points.
7. Press Enter. The logical volumes are listed by volume group.

   To determine the file systems, perform the following steps:

   a. Type `smitty` from your desktop window. The System Management Interface Tool is displayed.
   b. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel is displayed.
   c. Select **File Systems** and press Enter. The File Systems panel is displayed.
   d. Select **List All File Systems** to locate all file systems that have the same mount points as the logical volumes.
   e. Press Enter. The file systems are listed.
   f. Note the file system name of that volume group and the file system mount point, if it is mounted.
   g. Unmount these file systems.
8. Enter the following to convert the volume group from the original ESS hdisks to SDD vpaths:

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

When the conversion is complete, your application now accesses ESS 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 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. To see if SDD is installed, issue the following command from the root (/) directory:

  `lslpp -l *Sdd*`

  where 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.rte.432`

  ...

  ...

  An example of output from the `lslpp -l *Sdd*` or`lslpp -l ibmSdd_432.rte` command is:

```
Fileset                      Level    State        Description
----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  ibmSdd_432.rte             1.3.2.0  COMMITTED    IBM SDD
                                                   AIX V432 V433 for concurrent
                                                   HACMP

Path: /etc/objrepos
  ibmSdd_432.rte             1.3.2.0  COMMITTED    IBM SDD
                                                   AIX V432 V433 for concurrent
                                                   HACMP
```

- The ESS 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 is displayed.

    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.
  - Select **Devices** and press Enter. The Devices panel is displayed.
  - Select **Data Path Device** and press Enter. The Data Path Device panel is displayed.
  - Select **Display Data Path Device Configuration** and press Enter. A list is displayed of the pseudo devices and whether there are multiple paths configured for the devices.
- Make sure the SDD vpath devices you are going to migrate to do not belong to any other volume group, and that the corresponding physical device (ESS 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 there is no pv displayed for this vpath and its paths (hdisks). If a LUN has never belonged to any volume group, there is no pvid written on it. In case there is a

pvid written on the LUN and the LUN does not belong to any volume group, you need to clear the pvid from the LUN before using it to migrate a volume group. The commands to clear the pvid are:

```
chdev -l  hdiskN  -a  pv=clear
chdev -l  vpathN  -a  pv=clear
```

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

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

1. Add new SDD vpath devices to an existing non-SDD volume group:
   a. Type `smitty` and press Enter from your desktop window. The System Management Interface Tool panel is displayed.

      Tip:
      - The SMIT facility runs in two interfaces, nongraphical and graphical. This procedure uses the nongraphical interface. You can 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 is displayed.
   c. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel is displayed.
   d. Select **Volume Group** and press Enter. The Volume Group panel is displayed.
   e. Select **Add a Data Path Volume to a Volume Group** and press Enter.
   f. 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.
2. Mirror logical volumes from the original volume to a SDD ESS 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 is displayed. Remove all non-SDD devices.

Note: A non-SDD volume group can consist of non-ESS or ESS 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 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 is displayed.

      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 is displayed.
   c. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel is displayed.
   d. Select **Volume Group** and press Enter. The Volume Group panel is displayed.
   e. Select **Add a Data Path Volume to a Volume Group** and press Enter.
   f. 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 is displayed.

b. Select **System Storage Management (Physical & Logical)** and press Enter. The System Storage Management (Physical & Logical) panel is displayed.
c. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel is displayed.
d. Select **Volume Group** and press Enter. The Volume Group panel is displayed.
e. Select **Mirror a Volume Group** and press Enter. The Mirror a Volume Group panel is displayed.
f. 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 is displayed.
   b. Select **System Storage Management (Physical & Logical)** and press Enter. The System Storage Management (Physical & Logical) panel is displayed.
   c. Select **Logical Volume Manager** and press Enter. The Logical Volume Manager panel is displayed.
   d. Select **Volume Group** and press Enter. The Volume Group panel is displayed.
   e. Select **Synchronize LVM Mirrors** and press Enter. The Synchronize LVM Mirrors panel is displayed.
   f. Select **Synchronize by Physical Volume**.

   You can also 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 is displayed.
   b. Select **Logical Volumes** and press Enter. The Logical Volumes panel is displayed.
   c. Select **Set Characteristic of a Logical Volume** and press Enter. The Set Characteristic of a Logical Volume panel is displayed.
   d. Select **Remove Copy from a Logical Volume** and press Enter. The Remove Copy from a Logical Volume panel is displayed.

   You can also 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 is displayed.
   b. Select **Logical Volume manager** and press Enter. The Logical Volume manager panel is displayed.
   c. Select **Volume Groups** and press Enter. The Volume Groups panel is displayed.
   d. Select **Set Characteristics of a Volume Group** and press Enter. The Set Characteristics of a Volume Group panel is displayed.

e. Select **Remove a Physical Volume from a Volume Group** and press Enter. The Remove a Physical Volume from a Volume Group panel is displayed.

You can also 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. 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 state into the AIX errlog 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 8, “Using the datapath commands” on page 153.

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.

New and modified error log messages by SDD for HACMP

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. This SDD release is for HACMP environments only. See “SDD installation package attributes” on page 34 for more information on this release.

VPATH_DEVICE_OPEN

The 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. In addition, this error log message is posted when the vpath device fails to register its underlying paths or fails to read the persistent reserve key for the device.

VPATH_OUT_SERVICE

There is no path available to retry an I/O request that failed for a vpath device. The I/O request is returned to the calling program and this error log is posted.

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

Hardware

- ESS
- One or more pairs of fibre-channel host adapters

  To use SDD's input-output (I/O) load-balancing 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

Software

- 32-bit, Intel-based, Red Hat 7.2 Linux operating system running either:
  - Red Hat 2.4.9–31enterprise *vendor* kernel
- 32-bit, Intel-based, SuSE 7.3 Linux operating system running either:
  - SuSE k_smp-2.4.16-22 *vendor* kernel

Non-supported 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
- 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.

© Copyright IBM Corp. 1999, 2002

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 multi-path fibre-channel connection.

See the *IBM TotalStorage™ Enterprise Storage Server™: Introduction and Planning Guide* for more information about how to configure the ESS.

See the *IBM TotalStorage Enterprise Storage Server: Host Attachment Guide* for information on working around Linux LUN limitations.

Configuring fibre-channel adapters

You must configure the fibre-channel adapters and the adapters's drivers that are attached to your Linux host system before you install SDD. Follow the adapter-specific configuration instructions to configure the adapters.

See *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.

See *IBM TotalStorage Enterprise Storage Server: Host Attachment Guide* for working around Linux LUM 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 onto your host system as the root user.
2. Insert the SDD installation CD into your CD-ROM drive.
3. Type `mount /dev/cdrom /mnt/cdrom` to mount the CD-ROM drive.
4. Type `cd /mnt/cdrom` to change to the /mnt/cdrom directory.
5. If you're running Red Hat, type `cd redhat`, if you're running SuSE, type `cd suse`, 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 on the CD.

   (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 -iv IBMsdd-N.N.N.N-1.i686.rpm` to install 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 18 lists the SDD driver and its major component files.

Table 18. SDD components for a Linux host system

Filename	Location	Description
sdd-mod.o-xxxy	/opt/IBMsdd	SDD device driver file (where XXX stands for the kernel level of your host system and y represents smp or up
vpath.conf	/etc	SDD configuration file
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
```

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 to which your host system is attached and that the ESS 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, it is recommended that you 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 19.

Table 19. Summary of SDD commands for a Linux host system

Command	Description
cfgvpath	Configures vpath devices
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
sdd start	Loads the SDD driver and automatically configures disk devices for multi-path 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 multi-path access.

Using multiple commands for SDD configuration

Perform the following steps to manually load and configure SDD on your Linux host system:

1. Log onto 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 uni-processor 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 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.

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:

```
/dev/sda ( 8,  0) host=0 ch=0 id=0 lun=0 vid=COMPAQ  pid=BB01813467  serial=xxxxxxxx  X
/dev/sdb ( 8, 16) host=0 ch=0 id=1 lun=0 vid=COMPAQ  pid=BB01813467  serial=xxxxxxxx  X
/dev/sdc ( 8, 32) host=2 ch=0 id=0 lun=0 vid=IBM     pid=2105F20     serial=40418679
/dev/sdd ( 8, 48) host=2 ch=0 id=0 lun=1 vid=IBM     pid=2105F20     serial=40518679
/dev/sde ( 8, 64) host=2 ch=0 id=0 lun=2 vid=IBM     pid=2105F20     serial=40618679
/dev/sdf ( 8, 80) host=2 ch=0 id=0 lun=3 vid=IBM     pid=2105F20     serial=40718679
/dev/sdg ( 8, 96) host=2 ch=0 id=0 lun=4 vid=IBM     pid=2105F20     serial=50418679
/dev/sdh ( 8,112) host=2 ch=0 id=0 lun=5 vid=IBM     pid=2105F20     serial=50518679
/dev/sdi ( 8,128) host=2 ch=0 id=0 lun=6 vid=IBM     pid=2105F20     serial=50618679
/dev/sdj ( 8,144) host=2 ch=0 id=0 lun=7 vid=IBM     pid=2105F20     serial=50718679
```

   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 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 with `xxx` representing 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:

```
000 vpath0  (247,  0) 40418679 = /dev/sdc /dev/sdl /dev/sdcg /dev/sdcp
001 vpath1  (247, 16) 40518679 = /dev/sdd /dev/sdm /dev/sdch /dev/sdcq
002 vpath2  (247, 32) 40618679 = /dev/sde /dev/sdn /dev/sdci /dev/sdcr
003 vpath3  (247, 48) 40718679 = /dev/sdf /dev/sdo /dev/sdcj /dev/sdcs
004 vpath4  (247, 64) 50418679 = /dev/sdg /dev/sdp /dev/sdck /dev/sdct
005 vpath5  (247, 80) 50518679 = /dev/sdh /dev/sdq /dev/sdcl /dev/sdcu
006 vpath6  (247, 96) 50618679 = /dev/sdi /dev/sdr /dev/sdcm /dev/sdcv
007 vpath7  (247,112) 50718679 = /dev/sdj /dev/sds /dev/sdcn /dev/sdcw
008 vpath8  (247,128) 00818679 = /dev/sdk /dev/sdt /dev/sdco /dev/sdcx
```

   See Chapter 8, “Using the datapath commands” on page 153 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 74.

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 at the time 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 system startup:

1. Log onto 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)
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 system startup:

1. Log onto 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 HBA driver must already be loaded. This can be assured at start time by adding the appropriate driver(s) 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 command" on page 156 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 you desire to remain unconfigured. Removing the # allows previously excluded device to again be configured.

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
```

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 onto 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 -ql IBMsdd`.

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

```
package IBMsdd is not installed
```

Note: The `rmmod` command will not unload a driver that is in use.

Understanding the SDD error recovery policy

SDD, when in *multi-path mode*, makes it possible for you to concurrently download and install licensed internal code to the ESS while application IO 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 an ESS LUN. SDD in *multi-path* mode has the following characteristics:

- If an I/O error occurs on a path, SDD does not attempt to use the path again until 2000 successful I/O operations occur on an operational path. This process is known as *bypassing a path*. SDD bypasses a failing path twice (until the I/O error count reaches three) and then puts the path in the dead state.
- 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. SDD *never* puts the last operational path to an ESS LUN into the dead state even if I/O errors have occurred on the path.
- If an I/O error occurs on all the paths to an ESS LUN, SDD returns the I/O error back to the application.

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. It is possible to create a volume group or a vpath device with the *last single path*.

Note: SDD supports one SCSI adapter on the host system.

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

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

Partitioning SDD vpath devices

Disk partitions are known as *logical devices*. SDD for Linux only allows configuration of whole devices. The SDD naming scheme for disks and disk partitions follows the standard Linux disk naming convention. The following description illustrates the naming scheme for SCSI disks and disk partitions:

1. The first two letters indicate the SCSI device.
2. The next letter (or two letters), a-z, specifies the unique device name.
3. A number following the device name denotes the partition number. For example, /dev/sda is the whole device, while /dev/sda1 is a logical device representing the first partition of the whole device /dev/sda. Each device and partition has its own major and minor number.

Similarly then, a specific device file /dev/vpath*X* is created for each supported multi-path SCSI disk device (where *X* represents the unique device name; as with sd devices, *X* may be one or two letters).

Device files /dev/vpath*XY* are also created for each partition of the multi-path device (where *Y* represents the corresponding partition number). When a file system or user application wants to use the logical device, it should refer to /dev/vpath*XY* (for example, /dev/vpatha1 or /dev/vpathbc7) as its multi-path 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 4. SDD for a Windows NT host system

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

Hardware

- ESS
- Host system
- SCSI adapters and cables
- Fibre adapters and cables

Software

- Windows NT 4.0 operating system
- SCSI and fibre-channel device drivers

Non-support 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.
- 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 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.

© Copyright IBM Corp. 1999, 2002

SCSI requirements

To use the SDD SCSI support, ensure 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 requirements

To use the SDD fibre 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 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.

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 to your host system and required fibre-channel adapters 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, see *IBM Enterprise Storage Server Introduction and Planning 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. See *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.

Configuring SCSI adapters

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 the 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 97 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. Select the CD-ROM drive. A list of all the installed directories on the compact disc is displayed.
5. Select the \winNt\IBMsdd directory.
6. Run the setup.exe program. The Installshield starts.
7. Click **Next**. The Software License agreement is displayed.
8. Click **Yes**. 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 windows opens.
12. Select the type of setup you prefer from the following setup choices. IBM recommends that you select **Typical**.
    Typical
    Selects all options.
    Compact
    You select the minimum required options *only* (Subsystem Device Driver and the README file).
    Custom
    Select the options that you need.
13. Click **Next**. The Setup Complete window opens.

14. Click **Finish**. The SDD program prompts you to start your computer again.
15. 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 file

Note: You can use the **datapath query device** command to verify the SDD installation. SDD is successfully installed if the command runs successfully.

Configuring the 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. You must log on as an administrator user to have access to the Windows NT disk administrator.
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 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
```

4. Type `datapath query device` and press Enter. In the example shown in the following output, 10 devices are attached to the SCSI path:

```
Total Devices : 10

DEV#:   0  DEVICE NAME: Disk2 Part0  TYPE: 2105E20   SERIAL: 00A12028
============================================================================
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   SERIAL: 00A12028
============================================================================
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   SERIAL: 00B12028
============================================================================
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   SERIAL: 00B12028
============================================================================
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   SERIAL: 00D12028
============================================================================
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   SERIAL: 00D12028
============================================================================
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   SERIAL: 50812028
============================================================================
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   SERIAL: 50812028
============================================================================
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   SERIAL: 60012028
============================================================================
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   SERIAL: 60012028
============================================================================
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.
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” on page 86.

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 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   SERIAL: 00A12028
============================================================================
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   SERIAL: 00A12028
============================================================================
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   SERIAL: 00B12028
============================================================================
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   SERIAL: 00B12028
============================================================================
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   SERIAL: 00D12028
============================================================================
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   SERIAL: 00D12028
============================================================================
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   SERIAL: 50812028
============================================================================
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   SERIAL: 50812028
============================================================================
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   SERIAL: 60012028
============================================================================
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   SERIAL: 60012028
============================================================================
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's 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.

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 the Subsystem Device Driver” on page 92 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 loose access to your assigned volumes.
2. Install the new version of SDD. (See “Installing the SDD” on page 83 for instructions.)

Adding or modifying a multipath storage configuration to the ESS

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

Previous configuration with one additional path
```

3. Type `datapath query device` and press Enter. In the example shown in the following output, 10 devices are attached to the SCSI path:

```
Total Devices : 10

DEV#:   0  DEVICE NAME: Disk2 Part0  TYPE: 2105E20   SERIAL: 00A12028
===========================================================================
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   SERIAL: 00A12028
===========================================================================
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   SERIAL: 00B12028
===========================================================================
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   SERIAL: 00B12028
===========================================================================
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   SERIAL: 00D12028
===========================================================================
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   SERIAL: 00D12028
===========================================================================
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   SERIAL: 50812028
===========================================================================
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   SERIAL: 50812028
===========================================================================
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   SERIAL: 60012028
===========================================================================
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   SERIAL: 60012028
===========================================================================
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
```

Adding new storage to existing configuration

Perform the following steps to install additional storage:

1. Install any additional hardware to the ESS.
2. Configure the new storage to the server.
3. Restart the Windows NT server. Restarting will ensure correct multipath access to both existing and new storage and to your Windows NT server.

4. Verify that the new storage is added correctly. See “Verifying new storage is installed correctly”.

Verifying new storage is installed correctly

After adding new storage to existing configuration, you should verify the following conditions:

- The new storage is correctly installed and configured.
- The number of adapters and the number of paths to each ESS volume match the updated configuration.
- The Windows disk numbers of all primary paths are labeled as path #0.

Perform the following steps to verify that the additional storage has been installed correctly:

1. Click **Start —> Program —> Subsystem Device Driver —>Subsystem Device Driver Management**. An MS-DOS window opens.
2. 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
```

3. Type `datapath query device` and press Enter. The output includes information about any additional devices that were installed. In the example shown in the following output, the output includes information about the new devices that were assigned:

```
Total Devices : 16

DEV#:   0  DEVICE NAME: Disk2 Part0  TYPE: 2105E20   SERIAL: 00A12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk2 Part0     OPEN   NORMAL          9          0
    1   Scsi Port7 Bus0/Disk10 Part0     OPEN   NORMAL          5          0

DEV#:   1  DEVICE NAME: Disk2 Part1  TYPE: 2105E20   SERIAL: 00A12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk2 Part1     OPEN   NORMAL         26          0
    1   Scsi Port7 Bus0/Disk10 Part1     OPEN   NORMAL         38          0

DEV#:   2  DEVICE NAME: Disk3 Part0  TYPE: 2105E20   SERIAL: 00B12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk3 Part0     OPEN   NORMAL          9          0
    1   Scsi Port7 Bus0/Disk11 Part0     OPEN   NORMAL          7          0

DEV#:   3  DEVICE NAME: Disk3 Part1  TYPE: 2105E20   SERIAL: 00B12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk3 Part1     OPEN   NORMAL         34          0
    1   Scsi Port7 Bus0/Disk11 Part1     OPEN   NORMAL         30          0

DEV#:   4  DEVICE NAME: Disk4 Part0  TYPE: 2105E20   SERIAL: 31512028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk4 Part0     OPEN   NORMAL          8          0
    1   Scsi Port7 Bus0/Disk12 Part0     OPEN   NORMAL          6          0

DEV#:   5  DEVICE NAME: Disk4 Part1  TYPE: 2105E20   SERIAL: 31512028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk4 Part1     OPEN   NORMAL         35          0
    1   Scsi Port7 Bus0/Disk12 Part1     OPEN   NORMAL         28          0

DEV#:   6  DEVICE NAME: Disk5 Part0  TYPE: 2105E20   SERIAL: 00D12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk5 Part0     OPEN   NORMAL          5          0
    1   Scsi Port7 Bus0/Disk13 Part0     OPEN   NORMAL          9          0

DEV#:   7  DEVICE NAME: Disk5 Part1  TYPE: 2105E20   SERIAL: 00D12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk5 Part1     OPEN   NORMAL         28          0
    1   Scsi Port7 Bus0/Disk13 Part1     OPEN   NORMAL         36          0

DEV#:   8  DEVICE NAME: Disk6 Part0  TYPE: 2105E20   SERIAL: 40812028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk6 Part0     OPEN   NORMAL          5          0
    1   Scsi Port7 Bus0/Disk14 Part0     OPEN   NORMAL          9          0

DEV#:   9  DEVICE NAME: Disk6 Part1  TYPE: 2105E20   SERIAL: 40812028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk6 Part1     OPEN   NORMAL         25          0
    1   Scsi Port7 Bus0/Disk14 Part1     OPEN   NORMAL         38          0

DEV#:  10  DEVICE NAME: Disk7 Part0  TYPE: 2105E20   SERIAL: 50812028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk7 Part0     OPEN   NORMAL          7          0
    1   Scsi Port7 Bus0/Disk15 Part0     OPEN   NORMAL          7          0

DEV#:  11  DEVICE NAME: Disk7 Part1  TYPE: 2105E20   SERIAL: 50812028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk7 Part1     OPEN   NORMAL         34          0
    1   Scsi Port7 Bus0/Disk15 Part1     OPEN   NORMAL         30          0
```

```
DEV#:  12  DEVICE NAME: Disk8 Part0  TYPE: 2105E20   SERIAL: 60012028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk8 Part0     OPEN   NORMAL          7          0
    1   Scsi Port7 Bus0/Disk16 Part0     OPEN   NORMAL          7          0

DEV#:  13  DEVICE NAME: Disk8 Part1  TYPE: 2105E20   SERIAL: 60012028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk8 Part1     OPEN   NORMAL         29          0
    1   Scsi Port7 Bus0/Disk16 Part1     OPEN   NORMAL         35          0

DEV#:  14  DEVICE NAME: Disk9 Part0  TYPE: 2105E20   SERIAL: 00812028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk9 Part0     OPEN   NORMAL          6          0
    1   Scsi Port7 Bus0/Disk17 Part0     OPEN   NORMAL          8          0

DEV#:  15  DEVICE NAME: Disk9 Part1  TYPE: 2105E20   SERIAL: 00812028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port6 Bus0/Disk9 Part1     OPEN   NORMAL         28          0
    1   Scsi Port7 Bus0/Disk17 Part1     OPEN   NORMAL         36          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's 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 of the 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 from the Disk Administrator application.

Removing the Subsystem Device Driver

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 on the **Add/Remove** button.

Attention:

- After uninstalling the previous version, you must *immediately* install the new version of SDD to avoid any potential data loss. (See “Installing the SDD” on page 83 for instructions.)
- If you perform a system restart and accidentally overwrite the disk signature, you may permanently loose 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:

  support.microsoft.com

Displaying the current version of the Subsystem Device Driver

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 *your_installation_directory_letter:*\Winnt\system32\drivers directory where (*your_installation_directory_letter*\ is the directory letter where you have installed the sddpath.sys file).
3. Click the sddpath.sys file in *your_installation_directory_letter:*\Winnt\system32\drivers directory where *your_installation_directory_letter* refers to the letter of the directory in which you have installed the sddpath.sys file.
4. Right-click on the sddpath.sys file and then click **Properties**. The **sddpath.sys** properties window will open.
5. 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

SDD makes it possible for you to concurrently download and install the licensed internal code while applications continue running. SDD makes this process transparent to the Windows NT host system through its error recovery and retry algorithms. There are two modes of operation:

single-path mode

The 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 1000 maximum number of times to bypass the interval when the ESS host adapters are not available. This I/O error might be caused by the process of concurrently downloading licensed internal code. See “Concurrent download of ESS licensed internal code” on page 4 for more information.
- SDD never puts this single path into the DEAD state.

multiple-path mode

The Windows NT host system has multiple paths that are configured to an ESS LUN. SDD in multiple-path mode has the following characteristics:

- If an I/O error occurs on a path, SDD does not attempt to use the path again until 2000 successful I/O operations on an operational path. This process is known as *bypassing a path*. SDD bypasses a failing path twice (until the I/O error count reaches three) and then puts the path in the dead state.
- If I/O errors occur consecutively on a path and the I/O error count reaches three, SDD immediately puts the failing path into the dead state. It uses the same bypass algorithm twice on the failed path while it is in the dead state.
- SDD changes the failed path from the dead state back into the open state after a certain number of successful I/O operations on an operational path. The required number of successful I/O operations is 50 000 for Windows NT host system.
- If the first I/O operation fails after the path is put back into the open state, SDD puts the path into the dead state immediately and permanently. In this scenario, you must manually bring the path online by using the **datapath** command. However, 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. SDD *never* puts the last operational path to an ESS LUN into the dead state even if I/O errors have occurred on the path.
- If an I/O error occurs on all the paths to an ESS LUN, SDD returns the I/O error back to the application.

Using high availability clustering

The following items are required to support the Windows NT operating system in a clustering environment:

- SDD 1.2.1 or higher
- Windows NT 4.0 Enterprise Edition with service pack 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

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 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 the SDD” on page 83 section.
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 on page 96.

    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 on page 95 - 10 on page 95.

11. Install SDD on *server_2*, and restart *server_2*.

    For installation instructions, go to “Installing the SDD” on page 83 section.
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 hotfix Q305638 and restart *server_1*, reapply Service Pack 6A (or higher) to *server_1*, and restart *server_1* again.
    - Install the MSCS software on *server_2*. When server_2 is up, install hotfix Q305638 and restart *server_2*, reapply Service Pack 6A (or higher) 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 volumes and logical volumes for the host server. The secondary server only shows the physical volumes and the logical volumes that it owns.

Chapter 5. SDD for 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. 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 install SDD:

Hardware

- ESS
- Host system
- SCSI adapters and cables
- Fibre adapters and cables

Software

- Windows 2000 operating system with Service Pack 2 or higher
- Device driver for SCSI or fibre adapters

Non-supported environments

SDD does not support the following environments:

- A host system with both a SCSI channel and fibre-channel connection to a shared LUN.
- Windows 2000 operating system or a paging file on an SDD-controlled multipath 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.

ESS requirements

To successfully install SDD, make sure that you configure the ESS devices as IBM 2105*xxx* (where *xxx* is the ESS model number) on your Windows 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.

SCSI requirements

To use the SDD SCSI support, ensure 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.

© Copyright IBM Corp. 1999, 2002

- If you need the SDD I/O load-balancing and failover features ensure that a minimum of two SCSI adapters is 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 2000 host system, go to the following Web site:

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

Fibre requirements

To use the SDD fibre 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 is installed.

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

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 1.3.3.3 (or later) installation

Before installing SDD 1.3.3.3 (or later), you must:

Note: If you currently have SDD 1.3.2.x running, then IBM recommends an upgrade to 1.3.3.3 (or later). To upgrade to SDD 1.3.3.3 (or later), see "Upgrading the Subsystem Device Driver" on page 101.

1. Ensure that all hardware and software requirements are met before you install the Subsystem Device Driver. See "Verifying the hardware and software requirements" on page 97 for more information.
2. Configure the ESS to your host system. See "Configuring the ESS" for more information.
3. Configure the fibre-channel adapters that are attached to your Windows 2000 host system. See "Configuring fibre-channel adapters" on page 99 for more information.
4. Configure the SCSI adapters that are attached to your Windows 2000 host system. See "Configuring SCSI adapters" on page 99 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 the Subsystem Device Driver" on page 103
   b. "Installing SDD 1.3.3.3 (or later)" on page 100

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, see the *IBM Enterprise Storage Server Introduction and Planning 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.

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.

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*.

Planning for SDD 1.3.3.3 (or later) installation

The way you install SDD 1.3.3.3 (or later) depends on the SDD version you have currently installed on your Windows 2000 host system. There are four possible scenarios for installing SDD 1.3.3.3 (or later). The scenario you choose depends on the SDD version you have currently installed:

Note: It is required that Subsystem Device Driver 1.3.1.1 (or earlier) be uninstalled before Subsystem Device Driver 1.3.3.3 (or later) can be installed on your system.

Table 20. SDD 1.3.3.3 installation scenarios

Installation scenario	**SDD version you currently have installed**	**How to proceed**
Scenario 1	SDD 1.3.1.1 or earlier version	Go to: 1. "Uninstalling the Subsystem Device Driver" on page 103 2. "Installing SDD 1.3.3.3 (or later)" on page 100
Scenario 2	SDD 1.3.1.2 or SDD 1.3.2.x	Go to: 1. "Upgrading the Subsystem Device Driver" on page 101

Table 20. SDD 1.3.3.3 installation scenarios (continued)

Installation scenario	SDD version you currently have installed	How to proceed
Scenario 3	• SDD is currently installed in a cluster environment.	Go to: 1. "Upgrading to SDD 1.3.3.3 (or later) in a 2–node cluster environment" on page 102
Scenario 4	• SDD is not currently installed.	Perform the following steps: 1. Shut down your host system. 2. Disconnect all cables that connect the host bus adapters and the ESS devices. 3. "Installing SDD 1.3.3.3 (or later)"

Table 20 on page 99 further describes the installation scenarios and how you should proceed.

Installing SDD 1.3.3.3 (or later)

The following section describes how to install SDD 1.3.3.3 (or later) on your system.

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.
- If you have not already done so, disconnect all cables that connect the host bus adapters and the ESS devices (while your Windows 2000 host system is turned off) before proceeding with the installation procedures.

Perform the following steps to install SDD 1.3.3.3 (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. Select the CD-ROM drive. A list of all the installed directories on the compact disc is displayed.
5. Select the \win2k\IBMsdd directory (or your installation subdirectory)
6. Run the setup.exe program. The setup program starts.

   Tip:

   - If you have previously installed an earlier version of SDD, you will see an *"Upgrade?"* question while the setup program is running. You must answer n to this question. After exiting from the setup program, go to "Preparing for SDD 1.3.3.3 (or later) installation" on page 98 section for more information.
   - 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. Press Enter to exit the setup program once installation is complete.

   Note: You need to run **Start Type Configuration Utility** before shutting down your Windows 2000 host system if you're using dynamic disks.
8. Shut down your Windows 2000 host system.

9. Reconnect all cables that connect the host bus adapters and the ESS devices.
10. Restart your Windows 2000 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. Start Type Configuration Utility
4. README file

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 the SDD".
- Start Type Configuration Utility

Displaying the current version of the SDD

You can display the current version of SDD on a Windows 2000 host system by viewing the sddbus.sys file properties. Perform the following steps to view the properties of sddbus.sys file:

1. Click **Start —> Programs —> Accessories —> Windows Explorer** to open Windows Explorer.
2. In Windows Explorer, go to *your_installation_directory_drive_letter*:\Winnt\system32\drivers directory

   where *your_installation_directory_drive_letter* is the letter of the directory in which you have installed the sddbus.sys file.
3. Click the sddbus.sys file in *your_installation_directory_drive_leltter*:\Winnt\system32\drivers directory.
4. Right-click on the sddbus.sys file, and then click **Properties**. The sddbus.sys properties window opens.
5. In the sddbus.sys properties window, click **Version**. The file version and copyright information about the sddbus.sys file is displayed.

Upgrading the Subsystem Device Driver

IBM recommends that you perform the upgrade to SDD 1.3.3.3 (or later) if you currently have SDD 1.3.1.2 or 1.3.2.x installed on your Windows 2000 host system.

Perform the following steps for upgrading to SDD 1.3.3.3 (or later):

1. Shut down your Windows 2000 host system.
2. Disconnect all cables that connect the host bus adapters and the ESS devices.
3. Turn on your Windows 2000 host system.
4. Log on as the administrator user.
5. Insert the SDD installation CD-ROM into the selected drive.
6. Start the Windows 2000 Explorer program.
7. Select the CD-ROM drive. A list of all the installed directories on the compact disc is displayed.

8. Select the \win2k\IBMsdd directory (or your installation subdirectory)
9. Run the setup.exe program. The setup program starts.

   Note:

   - If you currently have SDD 1.3.1.1 (or earlier) installed on your Windows 2000 host system and *you did not perform an uninstallation* prior to installing 1.3.3.3 (or later), answer `n` to the *"Upgrade?"* question; After exiting from the setup program, go to Table 20 on page 99 for instructions on how you should proceed.
   - If you currently have SDD 1.3.1.2 or 1.3.2.0 installed on your Windows 2000 host system, answer `y` to the *"Upgrade?"* question.

10. Press Enter to exit the setup program once installation is complete.
11. Shut down your Windows 2000 host system.
12. Reconnect all cables that connect the host bus adapters and the ESS devices.
13. Turn on your Windows 2000 host system.

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 the SDD" on page 101.

Upgrading to SDD 1.3.3.3 (or later) in a 2–node cluster environment

The way you upgrade to SDD 1.3.3.3 (or later) in a 2-node cluster environment depends on what you have currently installed on your Windows 2000 host system.

Perform the following steps for upgrading to SDD 1.3.3.3 (or later) in a 2-node cluster environment if you have SDD 1.3.1.1 (or earlier) currently installed on your Windows 2000.

Scenario 1

You have SDD 1.3.1.1 (or earlier) currently installed on your Windows 2000 host system.

1. Move all cluster resources from node A to node B.
2. Follow the instructions from the "Uninstalling the Subsystem Device Driver" on page 103 section on node A.
3. Follow the instructions from the "Installing SDD 1.3.3.3 (or later)" on page 100 section.
4. When node A is up, move all resources from node B to node A.
5. Repeat steps 2 - 3 on node B.

Scenario 2

You have SDD 1.3.1.2 or SDD 1.3.2.x currently installed on your Windows 2000 host system.

1. Move all cluster resources from node A to node B.
2. Follow the instructions from the "Upgrading the Subsystem Device Driver" on page 101 section 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 the Subsystem Device Driver" on page 101 section on node B.

Uninstalling the Subsystem Device Driver

Perform the following steps to uninstall SDD on a Windows 2000 host system:

1. Shut down your Windows 2000 host system.
2. Disconnect all cables that connect the host bus adapters and the ESS devices.
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 on the **Add/Remove** button.
9. Restart your system.

   Attention:

   - You must install SDD 1.3.3.3 (or later) *immediately* after performing a system restart to avoid any potential data loss. Go to “Installing SDD 1.3.3.3 (or later)” on page 100 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 the Subsystem Device Driver in a 2–node cluster environment

IBM recommends the following steps if you intend to remove the multipathing functions to an ESS device in a 2–node cluster environment.

Perform the following steps to remove SDD 1.3.3.3 (or later) in a 2–node cluster environment:

1. Move all cluster resources from node A to node B.
2. Remove second HBA of node A from the switch configuration.
3. Remove shared access to the second HBA through the ESS Storwatch Specialist. Ensure LUNs are only accessed by the first HBA of node A.
4. Remove multiple ESS port access if applicable.
5. Uninstall SDD, see “Uninstalling the Subsystem Device Driver”.
6. Restart your system.
7. Move all cluster resources from node B to node A.
8. Perform steps 2 - 7 on node B.

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 Port1 Bus0  NORMAL   ACTIVE       4057          0      8       8
```

4. Type `datapath query device` and press Enter. In the example shown in the following output, 8 devices are attached to the SCSI path:

```
Total Devices : 8

DEV#:   0  DEVICE NAME: Disk7 Part7  TYPE: 2105E20   SERIAL: 01312028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk7 Part0     OPEN   NORMAL       1045          0

DEV#:   1  DEVICE NAME: Disk6 Part6  TYPE: 2105E20   SERIAL: 01212028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk6 Part0     OPEN   NORMAL        391          0

DEV#:   2  DEVICE NAME: Disk5 Part5  TYPE: 2105E20   SERIAL: 01112028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk5 Part0     OPEN   NORMAL       1121          0

DEV#:   3  DEVICE NAME: Disk4 Part4  TYPE: 2105E20   SERIAL: 01012028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk4 Part0     OPEN   NORMAL        332          0

DEV#:   4  DEVICE NAME: Disk3 Part3  TYPE: 2105E20   SERIAL: 00F12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk3 Part0     OPEN   NORMAL        375          0

DEV#:   5  DEVICE NAME: Disk2 Part2  TYPE: 2105E20   SERIAL: 31412028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk2 Part0     OPEN   NORMAL        258          0

DEV#:   6  DEVICE NAME: Disk1 Part1  TYPE: 2105E20   SERIAL: 31312028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk1 Part0     OPEN   NORMAL        267          0

DEV#:   7  DEVICE NAME: Disk0 Part0  TYPE: 2105E20   SERIAL: 31212028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk0 Part0     OPEN   NORMAL        268          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 on Disk drives.
5. Click on **Scan for hardware changes**.
6. 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 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. The following output is displayed:

```
Total Devices : 8

DEV#:   0  DEVICE NAME: Disk7 Part7  TYPE: 2105E20   SERIAL: 01312028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk7 Part0     OPEN   NORMAL        190          0
    1   Scsi Port2 Bus0/Disk15 Part0     OPEN   NORMAL        179          0

DEV#:   1  DEVICE NAME: Disk6 Part6  TYPE: 2105E20   SERIAL: 01212028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk6 Part0     OPEN   NORMAL        179          0
    1   Scsi Port2 Bus0/Disk14 Part0     OPEN   NORMAL        184          0

DEV#:   2  DEVICE NAME: Disk5 Part5  TYPE: 2105E20   SERIAL: 01112028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk5 Part0     OPEN   NORMAL        194          0
    1   Scsi Port2 Bus0/Disk13 Part0     OPEN   NORMAL        179          0

DEV#:   3  DEVICE NAME: Disk4 Part4  TYPE: 2105E20   SERIAL: 01012028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk4 Part0     OPEN   NORMAL        187          0
    1   Scsi Port2 Bus0/Disk12 Part0     OPEN   NORMAL        173          0

DEV#:   4  DEVICE NAME: Disk3 Part3  TYPE: 2105E20   SERIAL: 00F12028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk3 Part0     OPEN   NORMAL        215          0
    1   Scsi Port2 Bus0/Disk11 Part0     OPEN   NORMAL        216          0

DEV#:   5  DEVICE NAME: Disk2 Part2  TYPE: 2105E20   SERIAL: 31412028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk2 Part0     OPEN   NORMAL        115          0
    1   Scsi Port2 Bus0/Disk10 Part0     OPEN   NORMAL        130          0

DEV#:   6  DEVICE NAME: Disk1 Part1  TYPE: 2105E20   SERIAL: 31312028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk1 Part0     OPEN   NORMAL        122          0
    1    Scsi Port2 Bus0/Disk9 Part0     OPEN   NORMAL        123          0

DEV#:   7  DEVICE NAME: Disk0 Part0  TYPE: 2105E20   SERIAL: 31212028
============================================================================
Path#              Adapter/Hard Disk    State     Mode     Select     Errors
    0    Scsi Port1 Bus0/Disk0 Part0     OPEN   NORMAL        123          0
    1    Scsi Port2 Bus0/Disk8 Part0     OPEN   NORMAL        128          0
```

Error recovery and retry policy

SDD makes it possible for you to concurrently download and install the licensed internal code while applications continue running. SDD makes this process transparent to the Windows 2000 host system through its error recovery and retry algorithms. There are two modes of operation:

single-path mode

The 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 1000 maximum number of times to bypass the interval when the ESS host adapters are not available. This I/O error might be caused by the process of concurrently downloading licensed internal code. See “Concurrent download of ESS licensed internal code” on page 4 for more information.

- SDD never puts this single path into the DEAD state.

multiple-path mode

The Windows 2000 host system has multiple paths that are configured to an ESS LUN. SDD in multiple-path mode has the following characteristics:

- If an I/O error occurs on a path, SDD does not attempt to use the path again until 2000 successful I/O operations on an operational path. This process is known as *bypassing a path*. SDD bypasses a failing path twice (until the I/O error count reaches three) and then puts the path in the dead state.
- If I/O errors occur consecutively on a path and the I/O error count reaches three, SDD immediately puts the failing path into the dead state. It uses the same bypass algorithm twice on the failed path while it is in the dead state.
- SDD changes the failed path from the dead state back into the open state after a certain number of successful I/O operations on an operational path. The required number of successful I/O operations is 50 000 for Windows NT host system.
- If the first I/O operation fails after the path is put back into the open state, SDD puts the path into the dead state immediately and permanently. In this scenario, you must manually bring the path online by using the **datapath** command. However, 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. SDD *never* puts the last operational path to an ESS LUN into the dead state even if I/O errors have occurred on the path.
- If an I/O error occurs on all the paths to an ESS LUN, SDD returns the I/O error back to the application.

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, failover may 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

The following variables are used in this procedure:

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

Perform the following steps to configure a Windows 2000 cluster with SDD:

1. Configure LUNs on the 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 1.3.3.3 (or later)” on page 100 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 on page 108 - 10.

11. Install SDD on *server_2*, and restart *server_2*.

    For installation instructions, go to "Installing SDD 1.3.3.3 (or later)" on page 100 section.

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 only shows the physical volumes and the logical volumes that it owns.

Chapter 6. SDD for 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. For updated and additional information 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
- 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 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:

  www.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 which have been created and confirmed for multi-port access
- A fiber-optic cable to connect each fibre-channel adapter to an ESS port

Software

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

Non-supported 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.

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

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

Attention: For SDD to operate properly ensure that the patches in Table 22 on page 113 are installed on your HP-UX 11.0 host system.

© Copyright IBM Corp. 1999, 2002

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 to your host system and attach the required SCSI or fibre-channel adapters.

Configuring the ESS

Before you install SDD, configure your ESS for single-port or multiport access for each LUN. SDD requires a minimum of two independent paths that share the same logical unit to use the load-balancing and path-failover protection features.**With a single path, failover protection is not provided.**

For information about configuring your ESS, see *IBM Enterprise Storage Server Introduction and Planning 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 you have running. There are two types of special device files that 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 you have installed:

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

Table 21. SDD installation scenarios

Installation scenario	Description	How to proceed
Scenario 1	• SDD is not installed. • The SDD server for Expert is installed. • No software application or DBMS communicates directly to sdisk interface.	Go to: 1. “Determining if the SDD 1.3.1.5 (or later) server for Expert is installed” on page 113 2. “Installing the SDD” on page 114 3. “Standard UNIX applications” on page 124
Scenario 2	• SDD is not installed. • The SDD server for Expert is installed. • An existing application package or DBMS communicates directly to the sdisk interface.	Go to: 1. “Determining if the SDD 1.3.1.5 (or later) server for Expert is installed” on page 113 2. “Installing the SDD” on page 114 3. “Using applications with SDD” on page 124

Table 21. SDD installation scenarios (continued)

Scenario 3	• SDD is installed. • The SDD server for Expert is installed.	Go to "Determining if the SDD 1.3.1.5 (or later) server for Expert is installed""Upgrading the SDD" on page 117

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

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

Table 22. 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 & NMF, sync panics
32-bit	PHKL_20915	Trap-related panics/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, shmem, thread cumulative patch
32-bit or 64-bit	PHKL_21392	VxFS performance, hang, icache, DPFs
32-bit or 64-bit	PHKL_21624	start, JFS, PA8600, 3Gdata, NFS, IDS, PM, VM, async
32-bit or 64-bit	PHKL_21989	SCSI IO Subsystem Cumulative patch
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 SDDsrvHP*bb_yymmdd*.depot (where *bb* represents 32 or 64 bit, and *yymmdd* represents date of installation package). For ESS

Expert V2R1, the standalone SDD Server installation package is SDDsrvHP32_020115.depot for a 32-bit environment, and SDDsrvHP64_020115.depot for a 64-bit environment.)

- For instructions on how to remove the *standalone* version of the SDD server (for ESS Expert) from your HP 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 119

Installing the 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.
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 IBMdpo.depot file.

      For 32-bit mode applications, use:

      /*cdrom*/hp32bit/IBMdpo.depot

      or

      /*your_installation_directory*/hp32bit/IBMdpo.depot

      For 64-bit mode applications, use:

      /*cdrom*/hp64bit/IBMdpo.depot

      or

      /*your_installation_directory*/hp32bit/IBMdpo.depot
   c. Click **OK**.

      You will see output similar to that in either Figure 3 on page 115 or Figure 4 on page 115.

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

Figure 3. IBMdpo Driver 32-bit

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

Figure 4. IBMdpo Driver 64-bit

8. Click the **IBMdpo_tag** product.
9. Click **Actions** from the Bar menu, and then click **Mark for Install**.
10. Click **Actions** from the Bar menu, and then click **Install (analysis)**. An Install Analysis panel, is displayed showing the status of **Ready**.
11. Click **OK** to proceed. A Confirmation window opens which states that the installation will begin.
12. Type `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. Type `Yes` and press Enter. The installation of IBMdpo will now proceed.
14. An Install window opens informing you about the progress of the IBMdpo 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 completes.
15. Click **Done**. A Note window opens informing you that the local system will restart with the newly installed software.
16. Click **OK** to proceed. Before the machine console restarts, a message similar to the following is displayed:

    ```
    * A reboot of this system is being invoked. Please wait.

    *** FINAL System shutdown message (XXXXX) ***
    System going down IMMEDIATELY
    ```

Note: You can use the **datapath query device** command to verify the SDD installation. SDD is successfully installed if the command executes successfully.

Note: 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.

Post-installation

After SDD is installed, the device driver resides above the HP SCSI disk driver (sdisk) in the protocol stack. In other words, SDD now communicates to the HP-UX device layer. The SDD software installation procedure installs a number of SDD components and updates some system files. Those components and files are listed in the following tables.

Table 23. SDD components installed for HP host systems

File	Location	Description
libvpath.a	/usr/conf/lib	SDD device driver
vpath	/usr/conf/master.d	SDD configuration file
Executables	/opt/IBMdpo/bin	Configuration and status tools
README.sd	/opt/IBMdpo	README file
defvpath	/sbin	SDD configuration file used during startup
sddsrv	/sbin/sddsrv	SDD server daemon
sddsrv.conf	/etc/	SDD server configuration file

Table 24. System files updated for HP host systems

File	Location	Description
system	/stand/build	Forces the loading of the SDD device driver
lvmrc	/etc	Causes the **defvpath** command to run at start time

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

Command	Description
cfgvpath	Configures vpath devices
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 hdisks into SDD vpaths
vp2hd	Converts a volume group from SDD vpaths into ESS hdisks.

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 124 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 124 and read the information specific to the application you are using.

During the installation process, the following files were copied from the IBMdpo_depot to the system:

 # Kernel-related files

- /usr/conf/lib/libvpath.a
- /usr/conf/master.d/vpath

 # SDD driver related files

- /opt/IBMdpo
- /opt/IBMdpo/bin
- /opt/IBMdpo/README.sd
- /opt/IBMdpo/bin/cfgvpath
- /opt/IBMdpo/bin/datapath
- /opt/IBMdpo/bin/defvpath
- /opt/IBMdpo/bin/libvpath.a
- /opt/IBMdpo/bin/pathtest
- /opt/IBMdpo/bin/showvpath
- /opt/IBMdpo/bin/vpath
- /sbin/defvpath
- /sbin/sddsrv
- /etc/sddsrv.conf

In addition, the /stand/vmunix kernel was created with the device driver. The /stand/system directory was modified in order to add the device driver entry into the file. After these files were created, the /opt/IBMdpo/bin/cfgvpath program was initiated in order to create vpaths in the /dev/dsk and /dev/rdsk directories for all IBM disks which are available on the system. This information is stored in the /opt/IBMdpo file for use after restarting the machine.

Note: SDD devices are found in /dev/rdsk and /dev/dsk. The device is named according to the SDD number. A device with a number of 0 would be /dev/rdsk/vpath0.

Upgrading the SDD

Upgrading SDD consists of removing and reinstalling the IBMdpo package. If you are upgrading SDD, go to 117 and then go to “Installing the SDD” on page 114.

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. Restart or 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 **IBMdpo_tag** selection.
   a. Click **Actions** from the Bar menu, and then select **Mark for Remove**.
   b. Click **Actions** from the Bar menu, and then select **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. Type `Yes`. The analysis phase starts.
   e. After the analysis phase has finished, another Confirmation window opens indicating that the system will restarted after the uninstallation is complete. Type `Yes` and press Enter. The uninstallation of IBMdpo begins.
   f. An Uninstall window opens showing the progress of the IBMdpo 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   : IBMdpo_tag,..........
```

   The **Done** option is not available when the uninstall is in progress. It becomes available after the uninstall process completes.
9. Click **Done**. A Note window opens informing you that the local system will restart with the newly installed software.
10. Click **OK** to proceed. The following message is displayed on the machine console before it restarts:

```
* A reboot of this system is being invoked. Please wait.

*** FINAL System shutdown message (XXXXX) ***
System going down IMMEDIATELY
```

Note: 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 the SDD” on page 114.

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. Restart the system by typing:

   ```
   shutdown -r 0
   ```
2. Issue the **cfgvpath** command to reconfigure the vpath by typing:

   ```
   /opt/IBMdpo/bin/cfgvpath -c
   ```
3. Restart the system by typing:

   ```
   shutdown -r 0
   ```

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. The SDD server daemon starts automatically after the SDD driver package is installed. The SDD server daemon runs in the background at all times. The daemon scans to find failing paths (INVALID, CLOSE_DEAD, or DEAD) regularly between 2-5 minutes intervals. The daemon scans CLOSE paths and OPEN paths which are idle (failed paths) more often at 1 minute regular intervals.

Understanding how the SDD server daemon works

The SDD server 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 - 5 minute intervals between consecutive runs. If the test succeeds then SDD server reclaims these paths , and changes the states of these paths according to the following characteristics:

- If the state of the vpath is OPEN then the SDD server changes the state of INVALID and CLOSE_DEAD paths of that vpath to OPEN.
- If the state of the vpath is CLOSE then the SDD server changes the state of INVALID and CLOSE_DEAD paths of that vpath to CLOSE.
- The SDD server changes the state of DEAD paths to OPEN.

Path probing

The SDD server will regularly test CLOSE paths and OPEN paths which are idle to see if they are working or has become inoperational. The daemon sleeps for 1 minute intervals between consecutive runs. It will test OPEN and CLOSE paths, and detect if these paths have become idle. If the test fails, the SDD server then changes the state of these paths according to the following characteristics:

- If the vpath is in the OPEN state and the path is not working, then SDD server changes the state of the path from OPEN to DEAD.
- If vpath is in the CLOSE state and the path is not working, then SDD server changes the state of the path from CLOSE to CLOSE_DEAD.

Verifying if the SDD server has started

After you have installed SDD, verify if 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”.

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" for instructions.

Perform the following steps:

1. Type `nohup /sbin/sddsrv p`*`PN`* `&` 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" on page 119 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 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 support for single-path configuration

SDD 1.3.1.5 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.

Note: SDD supports one SCSI adapter on the host system.

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

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

The policy will be default to SINGLE_PATH_POLICY.

Understanding the SDD error recovery policy

The current SDD error recovery policy can cause a longer interruption for recovery to take place because of its delay in setting the status of broken paths to OFFLINE. The SDD server daemon corrects this delay because it regularly detects and recovers any broken paths. The SDD server then changes the status of the paths according on the its test results. See “SDD server daemon” on page 119 for more information.

The SDD error recovery policy has two modes of operation:

single-path mode

The 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 2 times to bypass the interval when the ESS host adapters are not available.
- SDD never puts this single path into the DEAD state.

  Note: SDD no longer supports single-path concurrent download of licensed internal code.

multiple-path mode

The host system has multiple paths that are configured to an ESS LUN. SDD in multiple-path mode has the following characteristics:

- If an I/O error occurs on a path, SDD does not attempt to use the path again until 2000 successful I/O operations on an operational path. This process is known as *bypassing a path*. SDD bypasses a failing path twice (until the I/O error count reaches three) and then puts the path in the dead state.
- If I/O errors occur consecutively on a path and the I/O error count reaches three, SDD immediately puts the failing path into the dead state. It uses the same bypass algorithm twice on the failed path while it is in the dead state.
- With auto path reclamation, the SDD server changes the status of a DEAD path to OPEN once it becomes operational. When the SDD server daemon detects a recovered path, the daemon changes the status of this path to ONLINE.
- The last path is reserved in OPEN status; Neither the SDD Subsystem Device Driver nor the SDD server daemon can change the last path to DEAD.
- 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 will be retried on all OPEN paths to an ESS LUN.

Converting a volume group

SDD provides two conversion scripts:

hd2vp script

The hd2vp script converts a volume group from ESS hdisks into SDD vpaths.

The syntax for hd2vp script is as follows:

```
hd2vp vgname
```

vp2hd The vp2hd script converts a volume group from SDD vpaths into ESS hdisks. Use the vp2hd program when you want to configure your applications back to original ESS hdisks, or when you want to remove the SDD from your HP host system.

The syntax for vp2hd script is as follows:

```
vp2hd vgname
```

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

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 other 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 123
- Create the volume group device directory. See “Creating the volume group device directory” on page 123
- Create the group special file. See “Creating the group special file” on page 123.

For more information about the **vgimport** command, see “Importing volume groups” on page 123

Exporting volume groups

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

–p The –p option previews the actions to be taken but do not update the /etx/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*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 node1.

Importing volume groups

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

–p The –p option previews the actions to be taken but do not update the /etx/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:

```
vgvimport –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 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 the SDD” on page 114. 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 are replaced 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 recreate 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

The task of creating a new logical volume to use SDD consists of the following subtasks.

Note: You must have super-user privileges to perform the following subtasks.

1. Determining the major number of the logical volume device
2. Creating a device node for the logical volume device
3. Creating a physical volume
4. Creating a volume group
5. Creating a logical volume
6. Creating a file system on the volume group
7. Mounting the logical volume.

To create a new logical volume that uses SDD, you first need to determine the major number of the logical volume device.

Determining 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              lvm
```

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

Creating a device node for the logical volume device: Creating a device node actually consists of:

1. Creating a subdirectory in the /dev directory for the volume group
2. Changing to the /dev directory
3. Creating a device node for the logical volume device

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 “Creating a physical volume” on page 126.

Creating a directory in /dev 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.

Next, change to the directory that you just created.

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

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

Next, recreate a device node for the logical volume device.

Recreating the physical volume: Type the following command to recreate the physical volume:

```
# pvcreate /dev/rdsk/vpath0
```

A message similar to the following is displayed:

```
Physical volume "/dev/rdsk/vpath0" has been successfully created.
```

In this example, the SDD device associated with the underlying disk is vpath0. Verify the underlying disk by typing the following **showvpath** command:

```
# /opt/IBMdpo/bin/showvpath
```

A message similar to the following is displayed:

```
vpath0:
 /dev/dsk/c3t4d0
```

Next, create the physical volume.

Creating a physical volume: Type the following command to create a physical volume:

```
# pvcreate /dev/rdsk/vpath0
```

Now create the volume group.

Creating a volume group: Type the following command to create a volume group:

```
# vgcreate /dev/vgibm /dev/dsk/vpath0
```

Now create the logical volume.

Creating a logical volume: Type the following command to create logical volume lvol1:

```
# lvcreate -L 100 -n lvol1 vgibm
```

The `-L 100` portion of the command makes a 100-MB volume group; you can make it larger if you want to. Now you are ready to create a file system on the volume group.

Creating a file system on the volume group: Type the following command to create a file system on the volume group:

```
# newfs -F hfs /dev/vgibm/rlvol1
```

Finally, mount the logical volume. This example assumes that you have a mount point called /mnt.

Mounting the logical volume: Type the following command to mount the logical volume lvol1:

```
# mount /dev/vgibm/lvol1 /mnt
```

Attention: In some cases it may be necessary to use standard HP 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 *HP-UX Reference Volume 2* at the following Web site:

docs.hp.com

Removing logical volumes

Perform the following procedures to remove logical volumes.

Removing the existing logical volume: Before the logical volume is removed, it must be unmounted. For example, type the following command to unmount logical volume lvol1:

```
# umount /dev/vgibm/lvol1
```

Next, remove the logical volume.

For example, type the following command to remove logical volume lvol1:

```
# lvremove /dev/vgibm/lvol1
```

A message similar to the following is displayed:

```
The logical volume "/dev/vgibm/lvol1" is not empty;
do you really want to delete the logical volume (y/n)
```

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

```
Logical volume "/dev/vgibm/lvol1" has been successfully removed.
Volume Group configuration for /dev/vgibm has been saved in
/etc/lvmconf/vgibm.conf
```

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

Next, remove the volume group.

Removing 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 recreate the logical volume.

Recreating the existing logical volumes

The task of converting an existing logical volume to use SDD consists of the following subtasks:

1. Determining the size of the logical volume
2. Recreating the physical volume
3. Recreating the volume group
4. Recreating the logical volume

5. Setting the correct timeout value for the logical volume manager

Note: You must have super-user privileges to perform these subtasks.

As an example, suppose you have a logical volume called lvol1 under a volume group vgibm, which is currently using the disk directly, (for example, through path /dev path /dev/dsk/c3t4d0). You would like to convert logical volume lvol1 to use SDD. To recreate the logical volume, you first need to determine the size of the logical volume.

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

```
# lvdisplay | 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. Next, recreate the physical volume.

Recreating the physical volume: Type the following command to recreate the physical volume:

```
# pvcreate /dev/rdsk/vpath0
```

A message similar to the following is displayed:

```
Physical volume "/dev/rdsk/vpath0" has been successfully created.
```

In this example, the SDD device associated with the underlying disk is vpath0. Verify the underlying disk by typing the following command:

```
# /opt/IBMdpo/bin/showvpath
```

A message similar to the following is displayed:

```
vpath0:
 /dev/dsk/c3t4d0
```

Next, recreate the volume group.

Recreating the volume group: Type the following command to recreate the volume group:

```
# vgcreate /dev/vgibm /dev/dsk/vpath0
```

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
```

Now recreate the logical volume.

Recreating the logical volume: Recreating the logical volume consists of a number of smaller steps:

1. Recreating the physical volume
2. Recreating the volume group
3. Recreating the logical volume
4. Setting the proper timeout value for the logical volume manager

Attention: The recreated logical volume should be the same size as the original volume; otherwise, the recreated volume cannot store the data that was on the original.

Type the following command to recreate 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.

Setting the correct 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:

1. Ensure that the timeout value for an SDD logical volume is set to the default. Type `lvdisplay /dev/vgibm/lvoly` and press Enter. If the timeout value is not default, type `lvchange -t 0 /dev/vgibm/lvoly` and press Enter to change it. (In this example, vgibm is the name of logical volume group that was previously configured to use SDD; in your environment the name may be different.)
2. 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/IBMdpo/bin/showvpath` and press Enter to obtain this information.
3. The recreated 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`, `vpchange`, or `vgreduce`, see the *HP-UX Reference Volume 2* at the following Web site:

docs.hp.com

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 the SDD” on page 114.
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/rdsk/vpathN
   ```

   In this example, *N* is the SDD device instance of the selected volume. Create mount points for the new file systems.
4. Install the file systems into the directory /etc/fstab. Click **yes** in the **mount at boot** field.
5. Install the file system mount points into the /etc/exports directory for export.
6. Restart the system.

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

Perform the following steps if you have the Network File System file server already configured to:

- Export file systems that reside on a multiport subsystem, and
- Use SDD partitions instead of sdisk partitions to access them

1. List the mount points for all currently exported file systems by looking in the /etc/exports directory.
2. Match the mount points found in step 1 with sdisk device link names (files named /dev/(r)dsk/cntndn) by looking in the /etc/fstab directory.
3. Match the sdisk device link names found in step 2 with SDD device link names (files named /dev/(r)dsk/vpathN) by issuing the **showvpath** command.
4. Make a backup copy of the current /etc/fstab file.
5. Edit the /etc/fstab file, replacing each instance of an sdisk device link named /dev/(r)dsk/cntndn with the corresponding SDD device link.
6. Restart the system.
7. Verify that each exported file system:
   - Passes the start time **fsck pass**
   - Mounts properly
   - Is exported and available to NFS clients

   If there is a problem with any exported file system after completing step 7, restore the original /etc/fstab file and restart to restore Network File System service. Then review your steps and try again.

Oracle

You must have super-user privileges to perform the following procedures. You also need to have Oracle documentation 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 in “Installing the SDD” on page 114.
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 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 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 the SDD” on page 114.
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 pre-installation tasks described in the *Oracle8 Installation Guide*. Plan the installation of Oracle8 on an file system residing on an SDD partition.
   b. Set up the Oracle user’s ORACLE_BASE and ORACLE_ HOME environment variables to the directories of this file system.
   c. Create two more SDD-resident file systems on two other SDD volumes. Each of the resulting three mount points should have a subdirectory named oradata. The subdirectory is used as a control file and redo log location for the installer’s default database (a sample database) as described in the *Oracle8 Installation Guide*. Oracle recommends using raw partitions for redo logs. To use SDD raw partitions as redo logs, create symbolic links from the three redo log locations to SDD raw device links (files named /dev/rdsk/vpathNs, where *N* is the SDD instance number, and *s* is the partition ID) that point to the slice.
3. Determine which SDD (vpathN) volumes you will use as Oracle8 database devices.
4. Partition the selected volumes using the HP-UX format utility. If SDD raw partitions are to be used by Oracle8 as database devices, be sure to leave disk cylinder 0 of the associated volume unused. This protects UNIX disk labels from corruption by Oracle8, as described in the *Oracle8 Installation Guide*.
5. Ensure that the Oracle software owner has read and write privileges to the selected SDD raw partition device files under the /devices directory.
6. Set up symbolic links from the oradata directory (under the first of the three mount points). Link the database files system <db>.dbf, tempdb.dbf, rbsdb.dbf, toolsd.bdbf, and usersdb.dbf to SDD raw device links (files named /dev/rdsk/vpathNs). Point to the partitions of the appropriate size, where *db* is the name of the database that you are creating. (The default is test.)

7. Install the Oracle8 server following the instructions in the *Oracle8 Installation Guide*. Be sure to be logged in as the Oracle software owner when you run the **orainst /m** command. Select the **Install New Product - Create Database Objects** option. Select **Raw Devices** for the storage type. Specify the raw device links set up in steps 2 and 6 for the redo logs and database files of the default database.
8. To set up other Oracle8 databases, you must set up control files, redo logs, and database files following the guidelines in the *Oracle8 Administrator's Reference*. Make sure any raw devices and file systems 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 the SDD" on page 114.
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.
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
   ```

   You now use the following SDD partition identifiers:

   ```
   mount /dev/dsk/vpath2 /oracle/mp1
   ```

   For example, assume that you found that vpath2 was the SDD identifier.

Follow the instructions in the *Oracle Installation Guide* for setting ownership and permissions.

If using raw partitions: Perform the following procedure if you have Oracle8 already installed and want to reconfigure it to use SDD partitions instead of sdisk partitions (for example, partitions accessed through /dev/rdsk/cntndn files).

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

- Changing 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.
- Recreating links to raw sdisk device links to point to raw SDD device links that access the same physical partitions.

Converting an Oracle installation from sdisk to SDD: Perform the following conversion steps:

1. Back up your Oracle8 database files, control files, and redo logs.
2. Obtain the sdisk device names for the Oracle8 mounted file systems by looking up the Oracle8 mount points in /etc/fstab and extracting the corresponding sdisk device link name (for example, /dev/rdsk/c1t4d0).
3. Launch the sqlplus utility.
4. 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/rdsk/c1tld0	oracle	dba	644	/dev/rdsk/vpath4

6. Fill in column 2 by issuing the command **ls -l** on each device link listed in column 1 and extracting the link source device file name.
7. Fill in the File Attributes columns by issuing the command **ls -l** on each Actual Device Node from column 2.
8. Install SDD following the instructions in the "Installing the SDD" on page 114.
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/IBMdpo/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, recreate 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 7. SDD for 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. 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
- One or more SCSI host adapters
- One or more fibre-channel host adapters. In case of a single-fibre 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 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.

  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

- A Sparc system running Solaris 2.6, Solaris 7, or 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

Nonsupported 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 a SDD pseudo device
- Root (/), /var, /usr, /opt, /tmp and swap partitions on a 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.

© Copyright IBM Corp. 1999, 2002

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 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, see *IBM Enterprise Storage Server Introduction and Planning 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.pkg** command looks similar to the following output:

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

Note:

- The installation package for the *standalone* version of the SDD server (for ESS Expert) is SDDsrvSUN*bb_yymmdd*.pkg. In this version*bb* represents 32 or 64 bit, and *yymmdd* represents the date of the installation package. For ESS Expert V2R1, the *standalone* SDD server installation package is SDDsrvSun32_020115.pkg for a 32-bit environment, and SDDsrvSun64_020115.pkg for a 64-bit environment.
- For instructions on how to remove the *standalone* version of the SDD server (for ESS Expert) from your 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 142

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 you install SDD depends on the kind of software 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's 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 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	• SDD is not installed. • No volume managers are installed. • No software application or DBMS is installed that communicates directly to the sd interface.	Go to: 1. "Installing the SDD" on page 138 2. "Standard UNIX applications" on page 145
Scenario 2	• SDD is not installed. • An existing volume manager, software application, or DBMS is installed that communicates directly to the sd interface.	Go to: 1. "Installing the SDD" on page 138 2. "Using applications with SDD" on page 145
Scenario 3	SDD is installed.	Go to: 1. "Upgrading the SDD" on page 141

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

Table 27. SDD package file names

Package file names	Description
sun32bit/IBMdpo	Solaris 2.6, Solaris 7, Solaris 8
sun64bit/IBMdpo	Solaris 7
sun64bit/IBMdpo	Solaris 8

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:

sunsolve.sun.com

For more information on the Solaris patches, see the *IBM TotalStorage Enterprise Storage Server: Host Attachment Guide*.

Attention: Analyze and study your operating system and application environment to ensure there are no conflicts with these patches prior to their installation.

Installing the 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.

Perform the following steps to install SDD on your Sun host system:

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 IBMdpo. For example,

   ```
   pkgadd -d /cdrom/cdrom0/sun32bit IBMdpo or
   pkgadd -d /cdrom/cdrom0/sun64bit IBMdpo
   ```

5. A message similar to the following message displays:

   ```
   Processing package instance <IBMdpo> from <var/spool/pkg>


   IBM DPO 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 <IBMdpo> [y,n,?]
   ```

6. Type `y` and press Enter to proceed. A message similar to the following message displays:

```
Installing IBM DPO driver as <IBMdpo>

## Installing part 1 of 1.
/etc/defvpath
/etc/rc2.d/S00vpath-config
/etc/rcS.d/S20vpath-config
/kernel/drv/vpathdd
/kernel/drv/vpathdd.conf
/opt/IBMdpo/cfgvpath
/opt/IBMdpo/datapath
/opt/IBMdpo/devlink.vpath.tab
/opt/IBMdpo/etc.system
/opt/IBMdpo/pathtest
/opt/IBMdpo/showvpath
/usr/sbin/vpathmkdev
[ verifying class <none>  ]
## Executing postinstall script.

DPO: Configuring 24 devices (3 disks * 8 slices)

Installation of <IBMdpo> was successful.

The following packages are available:
1 IBMcli ibm2105cli
         (sparc) 1.1.0.0
2 IBMdpo IBM DPO 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 displays:

```
*** 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.

DPO is now installed. Proceed to Post-Installation.
```

Verifying the SDD installation

To verify the SDD installation, perform the following steps:

1. Add `/opt/IBMdpo/bin` to the path.
   a. C shell: setenv PATH /opt/IBMdpo/bin:$PATH
   b. Bourne Shell: PATH=/opt/IBMdpo/bin:$PATH, export PATH
   c. Korn Shell: export PATH=/opt/IBMdpo/bin:$PATH

To verify that you successfully installed SDD, type `datapath query device`. If the command executes, SDD is installed.

Uninstalling an 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.

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. If you are using a database, restart the database.
4. Type `# pkgrm IBMdpo` and press Enter.

   Attention: A number of different installed packages display. Make sure you specify the correct package to uninstall.

   A message similar to the following message displays:

```
The following packages are available:
1 IBMcli ibm2105cli
         (sparc) 1.1.0.0
2 IBMdpo IBM DPO driver Version: May-10-2000 16:51
         (sparc) 1
```

5. Type `y` and press Enter. A message similar to the following message displays:

```
## Removing installed package instance <IBMdpo>

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
```

6. Type `y` and press Enter. A message similar to the following message displays:

```
## Verifying package dependencies.
## Processing package information.
## Executing preremove script.
Device busy
Cannot unload module: vpathdd
Will be unloaded upon reboot.
## Removing pathnames in class <none>
/usr/sbin/vpathmkdev
/opt/IBMdpo
/kernel/drv/vpathdd.conf
/kernel/drv/vpathdd
/etc/rcS.d/S20vpath-config
/etc/rc2.d/S00vpath-config
/etc/defvpath
## Updating system information.

Removal of <IBMdpo> 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 the SDD” on page 138 for detailed procedures.

Post-installation

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/rdsk and /dev/dsk directories. The device is named according to the SDD instance number. A device with an instance number of

0 would be: /dev/rdsk/vpath0a where *a* denotes the slice. Therefore, /dev/rdsk/vpath0c would be instance zero and slice 2.

After SDD is installed, the device driver resides above the Sun SCSI disk driver (sd) in the protocol stack. In other words, SDD now communicates to the Solaris device layer. The SDD software installation procedure installs a number of SDD components and updates some system files. Those components and files are listed in the following tables.

Table 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/IBMdpo/bin	Configuration and status tools
S20vpath-config	/etc/rcS.d	Boot initialization script **Note:** This script must come before other LVM initialization scripts
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	Configures vpath devices
showvpath	Lists all SDD devices and their underlying disks
vpathmkdev	Create SDD devices for /dev/dsk entries
datapath	SDD driver console command tool

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 145 and read the information specific to the application you are using.

Upgrading the SDD

Upgrading SDD consists of uninstalling and reinstalling the IBMdpo package. If you are upgrading SDD, go to “Uninstalling an SDD” on page 139 and then go to “Installing the SDD” on page 138.

Changing an SDD hardware configuration

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

1. Shut down the system. Type `shutdown -i0 -g0 -y` and press Enter.
2. Perform a configuration restart. From the OK prompt, type `boot -r` and press Enter. This uses the current SDD entries during restart, not the new entries. The restart forces the new disks to be recognized.
3. Run the SDD configuration utility to make the changes to the directory /opt/IBMdpo/bin. Type `cfgvpath -c` and press Enter.
4. Shut down the system. Type `shutdown -i6 -g0 -y` and press Enter.
5. After the restart, change to the /opt/IBMdpo/bin directory by typing:

   ```
   cd /opt/IBMdpo/bin
   ```
6. Type `devfsadm` and press Enter to reconfigure all the drives.
7. Type `vpathmkdev` and press Enter to create all the vpath devices.

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. The SDD server daemon starts automatically after the SDD driver package is installed. The SDD server daemon runs in the background at all times. The daemon scans to find failing paths (INVALID, CLOSE_DEAD, or DEAD) regularly between 2-5 minutes intervals. The daemon scans CLOSE paths and OPEN paths which are idle (failed paths) more often at 1 minute regular intervals.

Understanding how the SDD server daemon works

The SDD server 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 - 5 minute intervals between consecutive runs. If the test succeeds then SDD server reclaims these paths , and changes the states of these paths according to the following characteristics:

- If the state of the vpath is OPEN then the SDD server changes the state of INVALID and CLOSE_DEAD paths of that vpath to OPEN.
- If the state of the vpath is CLOSE then the SDD server changes the state of INVALID and CLOSE_DEAD paths of that vpath to CLOSE.
- The SDD server changes the state of DEAD paths to OPEN.

Path probing

The SDD server will regularly test CLOSE paths and OPEN paths which are idle to see if they are working or has become inoperational. The daemon sleeps for 1 minute intervals between consecutive runs. It will test OPEN and CLOSE paths, and detect if these paths have become idle. If the test fails, the SDD server then changes the state of these paths according to the following characteristics:

- If the vpath is in the OPEN state and the path is not working, then SDD server changes the state of the path from OPEN to DEAD.

- If vpath is in the CLOSE state and the path is not working, then SDD server changes the state of the path from CLOSE to CLOSE_DEAD.

Verifying if the SDD server has started

After you have installed SDD, verify if 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".

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 144 for instructions.

Perform the following steps:

1. Type `nohup /sbin/sddsrv p`*`PN`* `&` 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 p`*`22000`* `&`
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 other applications are not using.
   - If the SDD Server has started, the new port number that you specified is stored in the /etc/sddsrv.conf file and is used as a default port number until the next time that you change the port number.

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 by typing the following command: (where *pid* is the process id that you recorded in step 2)

   kill -9 pid

Understanding the SDD 1.3.1.0 (or later) support for single-path configuration

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.

Note: SDD supports one SCSI adapter on the host system.

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

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

Understanding the SDD 1.3.1.0 error recovery policy

The current SDD error recovery policy can cause a longer interruption for recovery to take place because of its delay in setting the status of broken paths to OFFLINE. The SDD server daemon corrects this delay because it regularly detects and recovers any broken paths. The SDD server then changes the status of the paths according to its test results. See “SDD server daemon” on page 142 for more information.

The SDD error recovery policy has two modes of operation:

single-path mode

The 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 2 times to bypass the interval when the ESS host adapters are not available.
- SDD never puts this single path into the DEAD state.

  Note: SDD no longer supports single-path concurrent download of licensed internal code.

multiple-path mode

The host system has multiple paths that are configured to an ESS LUN. SDD in multiple-path mode has the following characteristics:

- If an I/O error occurs on a path, SDD does not attempt to use the path again until 2000 successful I/O operations on an operational path. This

process is known as *bypassing a path*. SDD bypasses a failing path twice (until the I/O error count reaches three) and then puts the path in the dead state.

- If I/O errors occur consecutively on a path and the I/O error count reaches three, SDD immediately puts the failing path into the dead state. It uses the same bypass algorithm twice on the failed path while it is in the dead state.
- With auto path reclamation, the SDD server changes the state of a DEAD path to OPEN once it becomes operational. When the SDD server daemon detects a recovered path, the daemon changes the state of this path to ONLINE.
- The last path is reserved in OPEN state. Neither SDD nor the SDD server daemon can change the last path to DEAD.
- 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 will be retried on all OPEN paths to an ESS LUN.

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 the “Installing the SDD” on page 138 section. 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**, that 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 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 the SDD” on page 138 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/rdsk/vpathNs
   ```

   In this example, *N* is the SDD device instance of the selected volume. Create mount points for the new file systems.
5. Install the file systems into the /etc/fstab directory. Click `yes` in the **mount at boot** field.
6. Install the file system mount points into the directory /etc/exports for export.
7. Restart the system.

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

Perform the following steps if you have the Network File System file server already configured to:

- Export file systems that reside on a multiport subsystem
- To use SDD partitions instead of sd 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 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 exsists 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 superuser 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 in the “Installing the SDD” on page 138 section.
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 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 the SDD” on page 138 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 are files named /dev/rdsk/vpathNs, where*N* is the SDD instance number, and *s* is the partition ID.
3. Determine which SDD (vpathN) volumes you will use as Oracle8 database devices.
4. Partition the selected volumes using the Solaris format utility. If Oracle8 is to use SDD raw partitions as database devices, be sure to leave sector 0/disk cylinder 0 of the associated volume unused. This protects UNIX disk labels from corruption by Oracle8.
5. Ensure that the Oracle software owner has read and write privileges to the selected SDD raw partition device files under the /devices/pseudo directory.

6. Set up symbolic links in the oradata directory under the first of the three mount points. See step 2 on page 147. Link the database files to SDD raw device links (files named /dev/rdsk/vpathNs) that point to partitions of the appropriate size.
7. Install the Oracle8 server following the instructions in the *Oracle Installation Guide*. Be sure to be logged in as the Oracle software owner when you run the **orainst /m** command. Select the **Install New Product - Create Database Objects** option. Select **Raw Devices** for the storage type. Specify the raw device links set up in step 2 for the redo logs. Specify the raw device links set up in step 3 for the database files of the default database.
8. To set up other Oracle8 databases, you must set up control files, redo logs, and database files following the guidelines in the *Oracle8 Administrator's Reference*. Make sure any raw devices and file systems 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 the SDD" on page 138 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/clt8d2s2. (SDD partition identifiers end in an alphabetical character from a -g rather than s0, s1, s2, and so forth.) Write this down. A message similar to the following message displays:

```
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/rdsk/cntndn files).

All Oracle8 control, log, and data files are accessed either directly from mounted file systems or through links from the oradata subdirectory of each Oracle mount point set up on the server. Therefore, the process of converting an Oracle installation from sd to SDD is to change the Oracle mount points' physical devices in /etc/fstab from sd device partition links to the SDD device partition links that access the same physical partitions. After you do this, then re-create any links to raw sd device links to point to raw SDD device links that access the same physical partitions.

Converting an Oracle installation from sd to SDD partitions: Perform the following steps to convert an Oracle installation from sd to SDD partitions:

1. Back up your Oracle8 database files, control files, and redo logs.
2. Obtain the sd device names for the Oracle8 mounted file systems by looking up the Oracle8 mount points in /etc/vfstab and extracting the corresponding sd device link name (for example, /dev/rdsk/c1t4d0s4).
3. Launch the sqlplus utility.
4. 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 that each data file resides on. 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/rdsk/c1t1d0s4
   ```

   A message similar to the following message displays:

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

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

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

   A message similar to the following message displays:

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

7. Complete the basic installation steps in the "Installing the SDD" on page 138 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 the SDD” on page 138 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 partition of its own. 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 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 the SDD” on page 138 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 IBMdpo** 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 the SDD” on page 138 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 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 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 correctly set up 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 the SDD” on page 138 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/IBMdpo/bin/showvpath** command.
5. Unmount all current UFS logging file systems known to reside on the multiport subsystem through 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 8. 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	154
datapath query adaptstats	Displays performance information for all SCSI and FCS adapters that are attached to SDD devices	155
datapath query device	Displays information about devices	156
datapath query devstats	Displays performance information for a single SDD device or all SDD devices	158
datapath set adapter	Sets all device paths that are attached to an adapter to online or offline	160
datapath set path	Sets the path of a device to online or offline	161
datapath open path	Dynamically opens a path that is in an Invalid state.	162

© Copyright IBM Corp. 1999, 2002

datapath query adapter command

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.

datapath query adaptstats command

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/Write

datapath query device command

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

Parameters

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

Examples

If you type the `datapath query device 35` command, the following output is displayed:

```
DEV#:  35  DEVICE NAME: vpath0  TYPE: 2105E20   SERIAL: 60012028
===========================================================
Path#      Adapter/Hard Disk   State      Mode    Select        Errors
 0             scsi6/hdisk58    OPEN    NORMAL   7861147             0
 1             scsi5/hdisk36    OPEN    NORMAL   7762671             0
```

Note: 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.

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 that the path is attached to.

Hard Disk
The name of the logical device that the path is bound to.

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. It is 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 command

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

Syntax

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

Parameters

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

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

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/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 command

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 having 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 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

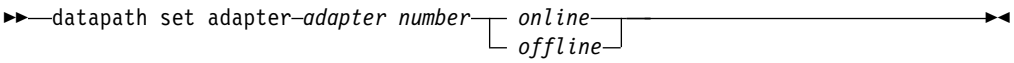

Parameters

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; while 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 command

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

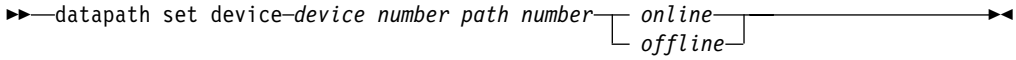

Parameters

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 command

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 AIX only.

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 is displayed:

```
Success: device 8 path 2 opened

DEV#:   8  DEVICE NAME: vpath9  TYPE: 2105E20   SERIAL: 20112028
POLICY:    Optimized
================================================================
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   SERIAL: 20112028
POLICY:    Optimized
================================================================
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 that the path is attached to

Hard Disk
The name of the logical device that the path is bound to

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 and output errors on this path.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS ″AS IS″ WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the 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.

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

© Copyright IBM Corp. 1999, 2002

DFSMS/MVS
ES/9000
ESCON
FICON
FlashCopy
HACMP/6000
IBM
Enterprise Storage Server
IBM TotalStorage
eServer
MVS/ESA
Netfinity
NetVista
NUMA-Q
Operating System/400
OS/390
OS/400
RS/6000
S/390
Seascape
SNAPSHOT
SP
StorWatch
System/360
System/370
System/390
TotalStorage
Versatile Storage Server
VM/ESA
VSE/ESA

Microsoft 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.

Linux is a trademark of Linus Torvalds and others.

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 erfuellt 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 purpose with regard to electromagnetic interference. If you find this is not suitable for your use, you may exchange it for a nonbusiness purpose one.

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 Agreement[1] 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.

1. 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 SCI.

This glossary uses the following cross-reference forms:

See This refers the reader to one of two kinds of related information:

- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
- A synonym or more preferred term

See also This refers the reader to one or more related terms.

Numerics

2105. The machine number for the IBM TotalStorage Enterprise Storage Server (ESS). Models of the ESS are expressed as the number 2105 followed by "Model <xxx>", such as 2105 Model 800. The 2105 Model 100 is an ESS expansion enclosure that is typically referred to simply as the Model 100. See also *IBM TotalStorage Enterprise Storage Server* and *Model 100*.

3390. The machine number of an IBM disk storage system. The ESS, when interfaced to IBM S/390 or zSeries hosts, is set up to appear as one or more 3390 devices, with a choice of 3390-2, 3390-3, or 3390-9 track formats.

3990. The machine number of an IBM control unit.

7133. The machine number of an IBM disk storage system. The Model D40 and 020 drawers of the 7133 can be installed in the 2105-100 expansion enclosure of the ESS.

8-pack. See *disk eight pack*.

A

access. (1) To obtain the use of a computer resource. (2) In computer security, a specific type of interaction between a subject and an object that results in flow of information from one to the other.

access-any mode. One of the two access modes that can be set for the ESS during initial configuration. It enables all fibre-channel-attached host systems with no defined access profile to access all logical volumes on the ESS. With a profile defined in ESS Specialist for a particular host, that host has access only to volumes that are assigned to the WWPN for that host. See also *pseudo-host* and *worldwide port name*.

ACK. See *request for acknowledgement and acknowledgement*.

active Copy Services server. The Copy Services server that manages the Copy Services domain. Either the primary or the backup Copy Services server can be the active Copy Services server. The backup Copy Services server is available to become the active Copy Services server if the primary Copy Services server fails. See also *backup Copy Services server*, *Copy Services client*, and *primary Copy Services server*.

alert. A message or log that a storage facility generates as the result of error event collection and analysis. An alert indicates that a service action is required.

allegiance. In Enterprise Systems Architecture/390, a relationship that is created between a device and one or more channel paths during the processing of certain conditions. See also *implicit allegiance*, *contingent allegiance*, and *reserved allegiance*.

© Copyright IBM Corp. 1999, 2002

allocated storage. In an ESS, the space that is allocated to volumes but not yet assigned. See also *assigned storage*.

American National Standards Institute (ANSI). An organization of producers, consumers, and general interest groups that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States. (A)

Anonymous. In ESS Specialist, the label on an icon that represents all connections that are using fibre-channel adapters between the ESS and hosts and that are not completely defined to the ESS. See also *anonymous host*, *pseudo-host*, and *access-any mode*.

anonymous host. Synonym for "pseudo-host" (in contrast to the Anonymous label that appears on some pseudo-host icons. See also *Anonymous* and *pseudo-host*.

ANSI. See *American National Standards Institute*.

APAR. See *authorized program analysis report*. (GC)

arbitrated loop. For fibre-channel connections, a topology that enables the interconnection of a set of nodes. See also *point-to-point connection* and *switched fabric*.

array. An ordered collection, or group, of physical devices (disk drive modules) that are used to define logical volumes or devices. More specifically, regarding the ESS, an array is a group of disks designated by the user to be managed by the RAID-5 technique. See also *redundant array of independent disks*.

ASCII. (American National Standard Code for Information Interchange) The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity check), that is used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A) Some organizations, including IBM, have used the parity bit to expand the basic code set.

assigned storage. On an ESS, the space allocated to a volume and assigned to a port.

authorized program analysis report (APAR). A report of a problem caused by a suspected defect in a current, unaltered release of a program. (GC)

availability. The degree to which a system or resource is capable of performing its normal function. See *data availability*.

B

backup Copy Services server. One of two Copy Services servers in a Copy Services domain. The other Copy Services server is the primary Copy Services server. The backup Copy Services server is available to become the active Copy Services server if the primary Copy Services server fails. A Copy Services server is software that runs in one of the two clusters of an ESS, and manages data-copy operations for that Copy Services server group. See also *active Copy Services server*, *Copy Services client*, and *primary Copy Services server*.

bay. In the ESS, the physical space used for installing SCSI, ESCON, and fibre-channel host adapter cards. The ESS has four bays, two in each cluster. See also *service boundary*.

bit. (1) Either of the digits 0 or 1 when used in the binary numeration system. (T) (2) The storage medium required to store a single binary digit. See also *byte*.

block. (1) A string of data elements recorded or transmitted as a unit. The elements may be characters, words, or physical records. (T) (2) In the ESS, a group of consecutive bytes used as the basic storage unit in fixed-block architecture (FBA). All blocks on the storage device are the same size (fixed size). See also *fixed-block architecture* and *data record*.

byte. (1) A group of eight adjacent binary digits that represent one EBCDIC character. (2) The storage medium required to store eight bits. See also *bit*.

C

cache. A special-purpose buffer storage, smaller and faster than main storage, used to hold a copy of instructions and data obtained from main storage and likely to be needed next by the processor. (T)

cache fast write. In the ESS, a form of the fast-write operation in which the storage server writes the data directly to cache, where it is available for later destaging.

cache hit. An event that occurs when a read operation is sent to the cluster, and the requested data is found in cache. The opposite of *cache miss*.

cache memory. Memory, typically volatile memory, that a storage server uses to improve access times to instructions or data. The cache memory is typically smaller and faster than the primary memory or storage medium. In addition to residing in cache memory, the same data also resides on the storage devices in the storage facility.

cache miss. An event that occurs when a read operation is sent to the cluster, but the data is not found in cache. The opposite of *cache hit*.

call home. A communication link established between the ESS and a service provider. The ESS can use this link to place a call to IBM or to another service provider

when it requires service. With access to the machine, service personnel can perform service tasks, such as viewing error logs and problem logs or initiating trace and dump retrievals. See also *heartbeat* and *remote technical assistance information network*.

cascading. (1) Connecting network controllers to each other in a succession of levels, to concentrate many more lines than a single level permits. (2) In high-availability cluster multiprocessing (HACMP), cascading pertains to a cluster configuration in which the cluster node with the highest priority for a particular resource acquires the resource if the primary node fails. The cluster node relinquishes the resource to the primary node upon reintegration of the primary node into the cluster.

catcher. A server that service personnel use to collect and retain status data that an ESS sends to it.

CCR. See *channel command retry*.

CCW. See *channel command word*.

CD. See *compact disc*.

compact disc. An optically read disc, typically storing approximately 660 MB. CD-ROM (compact disc read-only memory) refers to the read-only format used to distribute ESS code and documentation.

CEC. See *computer-electronic complex*.

channel. In Enterprise Systems Architecture/390, the part of a channel subsystem that manages a single I/O interface between a channel subsystem and a set of control units.

channel command retry (CCR). In Enterprise Systems Architecture/390, the protocol used between a channel and a control unit that enables the control unit to request that the channel reissue the current command.

channel command word (CCW). In Enterprise Systems Architecture/390, a data structure that specifies an I/O operation to the channel subsystem.

channel path. In Enterprise Systems Architecture/390, the interconnection between a channel and its associated control units.

channel subsystem. In Enterprise Systems Architecture/390, the part of a host computer that manages I/O communication between the program and any attached control units.

channel-subsystem image. In Enterprise Systems Architecture/390, the logical functions that a system requires to perform the function of a channel subsystem. With ESCON multiple image facility (EMIF), one channel subsystem image exists in the channel subsystem for each logical partition (LPAR). Each image appears to be an independent channel subsystem program, but all images share a common set of hardware facilities.

CKD. See *count key data*.

CLI. See *command-line interface*. See also *Copy Services command-line interface*.

cluster. (1) In the ESS, a partition capable of performing all ESS functions. With two clusters in the ESS, any operational cluster can take over the processing of a failing cluster. (2) In the AIX operating system, a group of nodes within a complex.

cluster processor complex (CPC). In the ESS, the unit within a cluster that provides the management function for the ESS. It consists of cluster processors, cluster memory, and related logic.

command-line interface (CLI). An interface provided by an operating system that defines a set of commands and enables a user (or a script-like language) to issue these commands by typing text in response to the command prompt (for example, DOS commands, UNIX shell commands). See also *Copy Services command-line interface*.

compression. (1) The process of eliminating gaps, empty fields, redundancies, and unnecessary data to shorten the length of records or blocks. (2) Any encoding that reduces the number of bits used to represent a given message or record. (GC)

computer-electronic complex (CEC). The set of hardware facilities associated with a host computer.

concurrent copy. A facility on a storage server that enables a program to make a backup of a data set while the logical volume remains available for subsequent processing. The data in the backup copy is frozen at the point in time that the server responds to the request.

concurrent installation of licensed internal code. Process of installing licensed internal code on an ESS while applications continue to run.

concurrent maintenance. Service that is performed on a unit while it is operational.

concurrent media maintenance. Service performed on a disk drive module (DDM) without losing access to the data.

configure. In storage, to define the logical and physical configuration of the input/output (I/O) subsystem through the user interface that the storage facility provides for this function.

consistent copy. A copy of a data entity (a logical volume, for example) that contains the contents of the entire data entity at a single instant in time.

console. A user interface to a server, such as can be provided by a personal computer. See also *IBM TotalStorage ESS Master Console*.

contingent allegiance. In Enterprise Systems Architecture/390, a relationship that is created in a control unit between a device and a channel when the channel accepts unit-check status. The allegiance causes the control unit to guarantee access; the control unit does not present the busy status to the device. The allegiance enables the channel to retrieve sense data that is associated with the unit-check status on the channel path associated with the allegiance.

control unit (CU). (1) A device that coordinates and controls the operation of one or more input/output devices, and synchronizes the operation of such devices with the operation of the system as a whole. (2) In Enterprise Systems Architecture/390, a storage server with ESCON, FICON, or OEMI interfaces. The control unit adapts a native device interface to an I/O interface supported by an ESA/390 host system. (3) In the ESS, the portion of the ESS that supports the attachment of emulated CKD devices over ESCON, FICON, or OEMI interfaces. See also *cluster*.

control-unit image. In Enterprise Systems Architecture/390, a logical subsystem that is accessed through an ESCON or FICON I/O interface. One or more control-unit images exist in each control unit. Each image appears as an independent control unit, but all control-unit images share a common set of hardware facilities. The ESS can emulate 3990-3, TPF, 3990-6, or 2105 control units.

control-unit initiated reconfiguration (CUIR). A software mechanism that the ESS uses to request that an operating system of an zSeries or S/390 host verify that one or more subsystem resources can be taken offline for service. The ESS can use this process to automatically vary channel paths offline and online to facilitate bay service or concurrent code installation. Depending on the operating system, support for this process might be model-dependent, might depend on the IBM TotalStorage Enterprise Storage Server Subsystem Device Driver, or might not exist.

Coordinated Universal Time (UTC). The international standard of time that is kept by atomic clocks around the world.

Copy Services client. Software that runs on each ESS cluster in the Copy Services server group and that performs the following functions:

- Communicates configuration, status, and connectivity information to the Copy Services server.
- Performs data-copy functions on behalf of the Copy Services server.

See also *active Copy Services server*, *backup Copy Services server*, and *primary Copy Services server*.

Copy Services CLI. See *Copy Services Command-Line Interface*.

Copy Services command-line interface (Copy Services CLI). In the ESS, command-line interface software provided with ESS Copy Services and used for invoking Copy Services functions from host systems attached to the ESS. See also *command-line interface*.

Copy Services server. An ESS cluster designated by the copy services administrator to perform the ESS Copy Services functions. See also *active Copy Services server*, *backup Copy Services server*, and *primary Copy Services server*.

Copy Services server group. A collection of user-designated ESS clusters participating in Copy Services functions managed by a designated active Copy Services server. A Copy Services server group is also called a Copy Services domain. See also *active Copy Services server*, *backup Copy Services server*, and *primary Copy Services server*.

count field. The first field of a count key data (CKD) record. This eight-byte field contains a four-byte track address (CCHH). It defines the cylinder and head that are associated with the track, and a one-byte record number (R) that identifies the record on the track. It defines a one-byte key length that specifies the length of the record's key field (0 means no key field). It defines a two-byte data length that specifies the length of the record's data field (0 means no data field). Only the end-of-file record has a data length of zero.

count key data (CKD). In Enterprise Systems Architecture/390, a data-record format employing self-defining record formats in which each record is represented by up to three fields: a *count* field identifying the record and specifying its format, an optional *key* field that can be used to identify the data area contents, and an optional *data* field that typically contains the user data. For CKD records on the ESS, the logical volume size is defined in terms of the device emulation mode (3390 or 3380 track format). The count field is always 8 bytes long and contains the lengths of the key and data fields, the key field has a length of 0 to 255 bytes, and the data field has a length of 0 to 65 535 or the maximum that will fit on the track. See also *data record*.

CPC. See *cluster processor complex*.

CRC. See *cyclic redundancy check*.

CU. See *control unit*.

CUIR. See *control-unit initiated reconfiguration*.

custom volume. In the ESS, a volume in count-key-data (CKD) format that is not a standard volume, which basically means that it does not necessarily present the same number of cylinders and capacity to its assigned logical control unit as provided

by one of the following standard S/390 volume types: 3390-2, 3390-3, 3390-9, 3390-2 (3380-track mode), or 3390-3 (3380-track mode). See also *count-key-data*, *interleave*, *standard volume*, and *volume*.

CUT. See *Coordinated Universal Time*.

cyclic redundancy check (CRC). A redundancy check in which the check key is generated by a cyclic algorithm. (T)

cylinder. A unit of storage on a CKD device with a fixed number of tracks.

D

DA. See *device adapter*. See also *SSA adapter*.

daisy chain. See *serial connection*.

DASD. See *direct access storage device*.

DASD fast write (DFW). A function of a storage server in which active write data is stored in nonvolatile cache, thus avoiding exposure to data loss.

data availability. The degree to which data is available when needed, typically measured as a percentage of time that the system would be capable of responding to any data request (for example, 99.999% available).

data compression. A technique or algorithm used to encode data such that the encoded result can be stored in less space than the original data. The original data can be recovered from the encoded result through a reverse technique or reverse algorithm. See also *compression*.

Data Facility Storage Management Subsystem. An operating environment that helps automate and centralize the management of storage. To manage storage, DFSMS provides the storage administrator with control over data class, storage class, management class, storage group, and automatic class selection routine definitions.

data field. The optional third field of a count key data (CKD) record. The count field specifies the length of the data field. The data field contains data that the program writes.

data record. The basic unit of S/390 and zSeries storage on an ESS, also known as a count-key-data (CKD) record. Data records are stored on a track. The records are sequentially numbered starting with 0. The first record, R0, is typically called the track descriptor record and contains data normally used by the operating system to manage the track. See also *count-key-data* and *fixed-block architecture*.

data sharing. The ability of multiple host systems to concurrently utilize data that they store on one or more storage devices. The storage facility enables configured storage to be accessible to any, or all, attached host systems. To use this capability, the host program must be designed to support data that it is sharing.

DDM. See *disk drive module*.

DDM group. See *disk eight pack*.

dedicated storage. Storage within a storage facility that is configured such that a single host system has exclusive access to the storage.

demote. To remove a logical data unit from cache memory. A storage server demotes a data unit to make room for other logical data units in the cache or because the logical data unit is not valid. The ESS must destage logical data units with active write units before they can be demoted.

destaging. Movement of data from an online or higher priority to an offline or lower priority device. The ESS stages incoming data into cache and then destages it to disk.

device. In Enterprise Systems Architecture/390, a disk drive.

device adapter (DA). A physical component of the ESS that provides communication between the clusters and the storage devices. The ESS has eight device adapters that it deploys in pairs, one from each cluster. DA pairing enables the ESS to access any disk drive from either of two paths, providing fault tolerance and enhanced availability.

device address. In Enterprise Systems Architecture/390, the field of an ESCON or FICON device-level frame that selects a specific device on a control-unit image.

device ID. In the ESS, the unique two-digit hexadecimal number that identifies the logical device.

device interface card. A physical subunit of a storage cluster that provides the communication with the attached DDMs.

device number. In Enterprise Systems Architecture/390, a four-hexadecimal-character identifier, for example 13A0, that the systems administrator associates with a device to facilitate communication between the program and the host operator. The device number is associated with a subchannel.

device sparing. A subsystem function that automatically copies data from a failing DDM to a spare DDM. The subsystem maintains data access during the process.

DFS. See *distributed file service*.

direct access storage device (DASD). (1) A mass storage medium on which a computer stores data. (2) A disk device.

disk cage. A container for disk drives. Each disk cage supports eight disk eight packs (64 disks).

disk drive. Standard term for a disk-based nonvolatile storage medium. The ESS uses hard disk drives as the primary nonvolatile storage media to store host data.

disk drive module (DDM). A field replaceable unit that consists of a single disk drive and its associated packaging.

disk drive module group. See *disk eight pack*.

disk eight pack. In the ESS, a group of eight disk drive modules (DDMs) installed as a unit in a DDM bay.

disk group. In the ESS, a collection of disk drives in the same SSA loop set up by the ESS to be available to be assigned as a RAID array. A disk group can be formatted as CKD or fixed block, and as RAID or non-RAID, or it can be left unformatted. A disk group is a logical assemblage of eight disk drives, in contrast to a disk eight pack. See also *disk eight pack*.

distributed file service (DFS). A service that provides data access over IP networks.

DNS. See *domain name system*.

domain. (1) That part of a computer network in which the data processing resources are under common control. (2) In TCP/IP, the naming system used in hierarchical networks. (3) A Copy Services server group, in other words, the set of clusters designated by the user to be managed by a particular Copy Services server.

domain name system (DNS). In TCP/IP, the server program that supplies name-to-address translation by mapping domain names to internet addresses. The address of a DNS server is the internet address of the server that hosts the DNS software for the network.

drawer. A unit that contains multiple DDMs and provides power, cooling, and related interconnection logic to make the DDMs accessible to attached host systems.

drive. (1) A peripheral device, especially one that has addressed storage media. See also *disk drive module*. (2) The mechanism used to seek, read, and write information on a storage medium.

duplex. (1) Regarding ESS Copy Services, the state of a volume pair after PPRC has completed the copy operation and the volume pair is synchronized. (2) In general, pertaining to a communication mode in which data can be sent and received at the same time.

dynamic sparing. The ability of a storage server to move data from a failing disk drive module (DDM) to a spare DDM while maintaining storage functions.

E

E10. The predecessor of the F10 model of the ESS. See also *F10*.

E20. The predecessor of the F20 model of the ESS. See also *F20*.

EBCDIC. See *extended binary-coded decimal interchange code*.

EC. See *engineering change*.

ECKD. See *extended count key data*.

eight pack. See *disk eight pack*.

electrostatic discharge (ESD). An undesirable discharge of static electricity that can damage equipment and degrade electrical circuitry.

emergency power off (EPO). A means of turning off power during an emergency, usually a switch.

EMIF. See *ESCON multiple image facility*.

enclosure. A unit that houses the components of a storage subsystem, such as a control unit, disk drives, and power source.

end of file. A coded character recorded on a data medium to indicate the end of the medium. On a CKD direct access storage device, the subsystem indicates the end of a file by including a record with a data length of zero.

engineering change (EC). An update to a machine, part, or program.

Enterprise Storage Server. See *IBM TotalStorage Enterprise Storage Server*.

Enterprise Systems Architecture/390 (ESA/390). An IBM architecture for mainframe computers and peripherals. Processor systems that follow the ESA/390 architecture include the ES/9000® family. See also *z/Architecture*.

Enterprise Systems Connection (ESCON). (1) An Enterprise Systems Architecture/390 and zSeries computer peripheral interface. The I/O interface uses ESA/390 logical protocols over a serial interface that configures attached units to a communication fabric. (2) A set of IBM products and services that provide a dynamically connected environment within an enterprise.

EPO. See *emergency power off*.

ERDS. See *error-recording data set*.

error-recording data set (ERDS). On S/390 and zSeries hosts, a data set that records data-storage and data-retrieval errors. A service information message (SIM) provides the error information for the ERDS.

ERP. See *error recovery procedure*.

error recovery procedure (ERP). Procedures designed to help isolate and, where possible, to recover from errors in equipment. The procedures are often used in conjunction with programs that record information on machine malfunctions.

ESA/390. See *Enterprise Systems Architecture/390*.

ESCD. See *ESCON director*.

ESCON. See *Enterprise System Connection*.

ESCON channel. An S/390 or zSeries channel that supports ESCON protocols.

ESCON director (ESCD). An I/O interface switch that provides for the interconnection of multiple ESCON interfaces in a distributed-star topology.

ESCON host systems. S/390 or zSeries hosts that attach to the ESS with an ESCON adapter. Such host systems run on operating systems that include MVS, VSE, TPF, or versions of VM.

ESCON multiple image facility (EMIF). In Enterprise Systems Architecture/390, a function that enables LPARs to share an ESCON channel path by providing each LPAR with its own channel-subsystem image.

EsconNet. In ESS Specialist, the label on a pseudo-host icon that represents a host connection that uses the ESCON protocol and that is not completely defined on the ESS. See also *pseudo-host* and *access-any mode*.

ESD. See *electrostatic discharge*.

eserver. See *IBM @server*.

ESS. See *IBM TotalStorage Enterprise Storage Server*.

ESS Copy Services. In the ESS, a collection of optional software features, with a Web-browser interface, used for configuring, managing, and monitoring data-copy functions.

ESS Copy Services CLI. See *Copy Services Command-Line Interface*.

ESS Expert. See *IBM TotalStorage Enterprise Storage Server Expert*.

ESS Master Console. See *IBM TotalStorage ESS Master Console*.

ESSNet. See *IBM TotalStorage Enterprise Storage Server Network*.

ESS Specialist. See *IBM TotalStorage Enterprise Storage Server Specialist*.

Expert. See *IBM TotalStorage Enterprise Storage Server Expert*.

extended binary-coded decimal interchange code (EBCDIC). A coding scheme developed by IBM used to represent various alphabetic, numeric, and special symbols with a coded character set of 256 eight-bit codes.

extended count key data (ECKD). An extension of the CKD architecture.

Extended Remote Copy (XRC). A function of a storage server that assists a control program to maintain a consistent copy of a logical volume on another storage facility. All modifications of the primary logical volume by any attached host are presented in order to a single host. The host then makes these modifications on the secondary logical volume.

extent. A continuous space on a disk that is occupied by or reserved for a particular data set, data space, or file. The unit of increment is a track. See also *multiple allegiance* and *parallel access volumes*.

F

F10. A model of the ESS featuring a single-phase power supply. It has fewer expansion capabilities than the Model F20.

F20. A model of the ESS featuring a three-phase power supply. It has more expansion capabilities than the Model F10, including the ability to support a separate expansion enclosure.

fabric. In fibre-channel technology, a routing structure, such as a switch, receives addressed information and routes to the appropriate destination. A fabric can consist of more than one switch. When multiple fibre-channel switches are interconnected, they are said to be *cascaded*.

failback. Cluster recovery from failover following repair. See also *failover*.

failover. In the ESS, the process of transferring all control of the ESS to a single cluster in the ESS when the other cluster in the ESS fails. See also *cluster*.

fast write. A write operation at cache speed that does not require immediate transfer of data to a disk drive. The subsystem writes the data directly to cache, to nonvolatile storage, or to both. The data is then available for destaging. A fast-write operation reduces the time an application must wait for the I/O operation to complete.

FBA. See *fixed-block architecture*.

FC. See *feature code*. **Note**: *FC* is a common abbreviation for fibre channel in the industry, but the ESS customer documentation library reserves *FC* for feature code.

FC-AL. See *Fibre Channel-Arbitrated Loop*.

FCP. See *fibre-channel protocol*.

FCS. See *fibre-channel standard*.

feature code (FC). A code that identifies a particular orderable option and that is used by service personnel to process hardware and software orders. Individual optional features are each identified by a unique feature code.

fibre channel. A data-transmission architecture based on the ANSI fibre-channel standard, which supports full-duplex communication. The ESS supports data transmission over fiber-optic cable through its fibre-channel adapters. See also *fibre-channel protocol* and *fibre-channel standard*.

Fibre Channel-Arbitrated Loop (FC-AL). An implementation of the fibre-channel standard that uses a ring topology for the communication fabric. Refer to American National Standards Institute (ANSI) X3T11/93-275. In this topology, two or more fibre-channel end points are interconnected through a looped interface. The ESS supports this topology.

fibre-channel connection (FICON). A fibre-channel communications protocol designed for IBM mainframe computers and peripherals.

fibre-channel protocol (FCP). A protocol used in fibre-channel communications with five layers that define how fibre-channel ports interact through their physical links to communicate with other ports.

fibre-channel standard (FCS). An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. The protocol has two layers. The IP layer defines basic interconnection protocols. The upper layer supports one or more logical protocols (for example, FCP for SCSI command protocols and SBCON for ESA/390 command protocols). Refer to American National Standards Institute (ANSI) X3.230-199x. See also *fibre-channel protocol*.

FICON. See *fibre-channel connection*.

FiconNet. In ESS Specialist, the label on a pseudo-host icon that represents a host connection that uses the FICON protocol and that is not completely defined on the ESS. See also *pseudo-host* and *access-any mode*.

field replaceable unit (FRU). An assembly that is replaced in its entirety when any one of its components fails. In some cases, a field replaceable unit might contain other field replaceable units. (GC)

FIFO. See *first-in-first-out*.

File Transfer Protocol (FTP). In TCP/IP, an application protocol used to transfer files to and from host computers. See also *Transmission Control Protocol/Internet Protocol*.

firewall. A protection against unauthorized connection to a computer or a data storage system. The protection is usually in the form of software on a gateway server that grants access to users who meet authorization criteria.

first-in-first-out (FIFO). A queuing technique in which the next item to be retrieved is the item that has been in the queue for the longest time. (A)

fixed-block architecture (FBA). An architecture for logical devices that specifies the format of and access mechanisms for the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size). The subsystem can access them independently.

fixed-block device. An architecture for logical devices that specifies the format of the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size); the subsystem can access them independently. This is the required format of the logical data units for host systems that attach with a SCSI or fibre-channel interface. See also *fibre-channel* and *small computer systems interface*.

FlashCopy. An optional feature for the ESS that can make an instant copy of data, that is, a point-in-time copy of a volume.

FRU. See *field replaceable unit*.

FTP. See *File Transfer Protocol*.

full duplex. See *duplex*.

G

GB. See *gigabyte*.

GDPS. See *Geographically Dispersed Parallel Sysplex*.

Geographically Dispersed Parallel Sysplex (GDPS). An S/390 multisite application- availability solution.

gigabyte (GB). A gigabyte of storage is 10^9 bytes. A gigabyte of memory is 2^{30} bytes.

group. In ESS documentation, a nickname for two different kinds of groups, depending on the context. See *disk eight pack* or *Copy Services server group*.

H

HA. See *host adapter*.

HACMP. See *High-Availability Cluster Multi-Processing*.

hard disk drive (HDD). (1) A storage medium within a storage server used to maintain information that the storage server requires. (2) A mass storage medium for computers that is typically available as a fixed disk (such as the disks used in system units of personal computers or in drives that are external to a personal computer) or a removable cartridge.

hardware service manager (HSM). An option on an AS/400 or iSeries host that enables the user to display and work with system hardware resources and to debug input-output processors (IOP), input-output adapters (IOA), and devices.

HDA. See *head and disk assembly*.

HDD. See *hard disk drive*.

hdisk. An AIX term for storage space.

head and disk assembly (HDA). The portion of an HDD associated with the medium and the read/write head.

heartbeat. A status report sent at regular intervals from the ESS. The service provider uses this report to monitor the health of the call home process. See also *call home*, *heartbeat call home record*, and *remote technical assistance information network*.

heartbeat call home record. Machine operating and service information sent to a service machine. These records might include such information as feature code information and product logical configuration information.

hierarchical storage management. (1) A function provided by storage management software such as Tivoli Storage Management or Data Facility Storage Management Subsystem/MVS (DFSMS/MVS) to automatically manage free space based on the policy that the storage administrator sets. (2) In AS/400 storage management, an automatic method to manage and distribute data between the different storage layers, such as disk units and tape library devices.

High-Availability Cluster Multi-Processing (HACMP). Software that provides host clustering, so that a failure of one host is recovered by moving jobs to other hosts within the cluster.

high-speed link (HSL). A hardware connectivity architecture that links system processors to system input/output buses and other system units.

home address (HA). A nine-byte field at the beginning of a track that contains information that identifies the physical track and its association with a cylinder.

Note: In the ESS, the acronym HA is shared between home address and host adapter. See also *host adapter*.

hop. Interswitch connection. A hop count is the number of connections that a particular block of data traverses between source and destination. For example, data traveling from one hub over a wire to another hub traverses one hop.

host. See *host system*.

host adapter (HA). A physical subunit of a storage server that provides the ability to attach to one or more host I/O interfaces. The Enterprise Storage Server has four HA bays, two in each cluster. Each bay supports up to four host adapters.

In the ESS, the acronym HA is shared between home address and host adapter. See also *home address*.

host name. The Internet address of a machine in the network. In the ESS, the host name can be entered in the host definition as the fully qualified domain name of the attached host system, such as `mycomputer.city.company.com`, or as the subname of the fully qualified domain name, for example, `mycomputer`. See also *host system*.

host processor. A processor that controls all or part of a user application network. In a network, the processing unit in which the data communication access method resides. See also *host system*.

host system. A computer, either of the mainframe (S/390 or zSeries) or of the open-systems type, that is connected to the ESS. S/390 or zSeries hosts are connected to the ESS through ESCON or FICON interfaces. Open-systems hosts are connected to the ESS by SCSI or fibre-channel interfaces.

hot plug. Pertaining to the ability to add or remove a hardware facility or resource to a unit while power is on.

HSL. See *high-speed link*.

HSM. See *hierarchical storage management* or *Hardware Service Manager*.

I

IBM @server. The IBM brand name for a series of server products that are optimized for e-commerce. The products include the iSeries, pSeries, xSeries, and zSeries.

IBM product engineering (PE). The third-level of IBM service support. Product engineering is composed of IBM engineers who have experience in supporting a product or who are knowledgeable about the product.

IBM TotalStorage. The brand name used to identify storage products from IBM, including the IBM TotalStorage Enterprise Storage Server (ESS). See also *IBM TotalStorage Enterprise Storage Server* and *IBM TotalStorage Enterprise Storage Server Specialist*.

IBM TotalStorage Enterprise Storage Server (ESS). A member of the Seascape product family of storage servers and attached storage devices (disk drive modules). The ESS provides for high-performance, fault-tolerant storage and management of enterprise data, providing access through multiple concurrent operating systems and communication protocols. High performance is provided by multiple symmetric multiprocessors, integrated caching, RAID support for the disk drive modules, and disk access through a high-speed serial storage architecture (SSA) interface.

IBM TotalStorage Enterprise Storage Server Expert (ESS Expert). 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-striping methodology.

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

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

NVS. See *nonvolatile storage*.

O

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

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

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

organizationally unique identifier (OUI). An IEEE-standards number that identifies an organization with a 24-bit globally unique assigned number referenced by various standards. OUI is used in the family of 802 LAN standards, such as Ethernet and Token Ring.

original equipment manufacturer's information (OEMI). A reference to an IBM guideline for a computer peripheral interface. The interface uses ESA/390 logical protocols over an I/O interface that configures attached units in a multidrop bus topology.

OS/390. The IBM operating system that includes and integrates functions previously provided by many IBM software products (including the MVS operating system) for the IBM S/390 family of enterprise servers.

OS/400. The IBM operating system that runs the IBM AS/400 and iSeries @server families of servers.

OUI. See *organizationally unique identifier*.

P

panel. The formatted display of information that appears on a display screen.

parallel access volume (PAV). An advanced function of the ESS that enables OS/390 and z/OS systems to issue concurrent I/O requests against a CKD logical volume by associating multiple devices of a single control-unit image with a single logical device. Up to eight device addresses can be assigned to a PAV. The PAV function enables two or more concurrent write operations to the same logical volume, as long as the write operations are not to the same extents. See also *extent*, *I/O Priority Queueing*, and *multiple allegiance*.

parity. A data checking scheme used in a computer system to ensure the integrity of the data. The RAID implementation uses parity to re-create data if a disk drive fails.

path group. In ESA/390 architecture, a set of channel paths that are defined to a control unit as being associated with a single logical partition (LPAR). The channel paths are in a group state and are online to the host. See also *logical partition*.

path group identifier. In ESA/390 architecture, the identifier that uniquely identifies a given logical partition (LPAR). The path group identifier is used in communication between the LPAR program and a device. The identifier associates the path group with one or more channel paths, thereby defining these paths to the control unit as being associated with the same LPAR. See also *logical partition*.

PAV. See *parallel access volume*.

PCI. See *peripheral component interconnect*.

PE. See *IBM product engineering*.

Peer-to-Peer Remote Copy (PPRC). A function of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also *synchronous PPRC* and *PPRC Extended Distance*.

peripheral component interconnect (PCI). An architecture for a system bus and associated protocols that supports attachments of adapter cards to a system backplane.

physical path. A single path through the I/O interconnection fabric that attaches two units. For Copy Services, this is the path from a host adapter on one ESS (through cabling and switches) to a host adapter on another ESS.

point-to-point connection. For fibre-channel connections, a topology that enables the direct interconnection of ports. See *arbitrated loop* and *switched fabric*.

port. In the ESS, a physical connection on a host adapter to the cable that connects the ESS to hosts, switches, or another ESS. The ESS uses SCSI and ESCON host adapters that have two ports per adapter, and fibre-channel host adapters that have one port. See also *ESCON*, *fibre channel*, *host adapter*, and *small computer system interface*.

POST. See *power-on self test*.

power-on self test (POST). A diagnostic test that servers or computers run when they are turned on.

PPRC. See *Peer-to-Peer Remote Copy*.

PPRC Extended Distance. An optional feature for the ESS that maintains a fuzzy copy of a logical volume on the same ESS or on another ESS. In other words, all modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume at a later point in time. The original order of update is not strictly maintained. See also *Peer-to-Peer Remote Copy (PPRC)* and *synchronous PPRC*.

PPRC-XD. See *PPRC Extended Distance*.

predictable write. A write operation that can cache without knowledge of the existing format on the medium. All write operations on FBA DASD devices are predictable. On CKD DASD devices, a write operation is predictable if it does a format write operation for the first data record on the track.

primary Copy Services server. One of two Copy Services servers in a Copy Services server group. The primary Copy Services server is the active Copy Services server until it fails; it is then replaced by the backup Copy Services server. A Copy Services server is software that runs in one of the two clusters of an ESS and performs data-copy operations within that group. See *active Copy Services server* and *backup Copy Services server*.

product engineering. See *IBM product engineering*.

program. On a computer, a generic term for software that controls the operation of the computer. Typically, the program is a logical assemblage of software modules that perform multiple related tasks.

program-controlled interruption. An interruption that occurs when an I/O channel fetches a channel command word with the program-controlled interruption flag on.

program temporary fix (PTF). A temporary solution or bypass of a problem diagnosed by IBM in a current unaltered release of a program. (GC)

promote. To add a logical data unit to cache memory.

protected volume. In 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*.

Seascape architecture. A storage system architecture developed by IBM for open-systems servers and S/390 and zSeries host systems. It provides storage solutions that integrate software, storage management, and technology for disk, tape, and optical storage.

self-timed interface (STI). An interface that has one or more conductors that transmit information serially between two interconnected units without requiring any clock signals to recover the data. The interface performs clock recovery independently on each serial data stream and uses information in the data stream to determine character boundaries and inter-conductor synchronization.

sequential access. A mode of accessing data on a medium in a manner that requires the storage device to access consecutive storage locations on the medium.

sequential access method (SAM). An access method for storing, deleting, or retrieving data in a continuous sequence based on the logical order of the records in the file.

serial connection. A method of device interconnection for determining interrupt priority by connecting the interrupt sources serially.

serial storage architecture (SSA). An IBM standard for a computer peripheral interface. The interface uses a SCSI logical protocol over a serial interface that configures attached targets and initiators in a ring topology. See also *SSA adapter*.

server. (1) A host that provides certain services to other hosts that are referred to as clients. (2) A functional unit that provides services to one or more clients over a network. (GC)

service boundary. A category that identifies a group of components that are unavailable for use when one of the components of the group is being serviced. Service boundaries are provided on the ESS, for example, in each host bay and in each cluster.

service information message (SIM). A message sent by a storage server to service personnel through an S/390 operating system.

service personnel. A generalization referring to individuals or companies authorized to service the ESS. The terms “service provider”, “service representative”, and “IBM service support representative (SSR)” refer to types of service personnel. See also *service support representative*.

service processor. A dedicated processing unit used to service a storage facility.

service support representative (SSR). Individuals or a company authorized to service the ESS. This term also refers to a service provider, a service representative, or an IBM service support representative (SSR). An IBM SSR installs the ESS.

shared storage. In an ESS, storage that is configured so that multiple hosts can concurrently access the storage. The storage has a uniform appearance to all hosts. The host programs that access the storage must have a common model for the information on a storage device. The programs must be designed to handle the effects of concurrent access.

shortwave laser adapter. A connector used between host and ESS to support shortwave fibre-channel communication.

SIM. See *service information message*.

Simple Network Management Protocol (SNMP). In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application layer protocol. Information on devices managed is defined and stored in the application’s Management Information Base (MIB). (GC) See also *management information base*.

simplex volume. A volume that is not part of a FlashCopy, XRC, or PPRC volume pair.

Single-Byte Command Code Sets Connection (SBCON). The ANSI standard for the ESCON or FICON I/O interface.

small computer system interface (SCSI). A standard hardware interface that enables a variety of peripheral devices to communicate with one another. (GC)

smart relay host. A mail relay or mail gateway that has the capability to correct e-mail addressing problems.

SMIT. See *System Management Interface Tool*.

SMP. See *symmetric multi-processor*.

SNMP. See *simple network management protocol*.

software transparency. Criteria applied to a processing environment that states that changes do not require modifications to the host software in order to continue to provide an existing function.

spare. A disk drive on the ESS that can replace a failed disk drive. A spare can be predesignated to allow automatic dynamic sparing. Any data preexisting on a disk drive that is invoked as a spare is destroyed by the dynamic sparing copy process.

spatial reuse. A feature of serial storage architecture that enables a device adapter loop to support many simultaneous read/write operations. See also *serial storage architecture*.

Specialist. See *IBM TotalStorage Enterprise Storage Server Specialist*.

SSA. See *serial storage architecture*.

SSA adapter. A physical adapter based on serial storage architecture. SSA adapters connect disk drive modules to ESS clusters. See also *serial storage architecture*.

SSID. See *subsystem identifier*.

SSR. See *service support representative*.

stacked status. In Enterprise Systems Architecture/390, the condition when the control unit is in a holding status for the channel, and the last time the control unit attempted to present the status, the channel responded with the stack-status control.

stage operation. The operation of reading data from the physical disk drive into the cache.

staging. To move data from an offline or low-priority device back to an online or higher priority device, usually on demand of the system or on request of the user.

standard volume. In the ESS, a volume that emulates one of several S/390 volume types, including 3390-2, 3390-3, 3390-9, 3390-2 (3380-track mode), or 3390-3 (3380-track mode), by presenting the same number of cylinders and capacity to the host as provided by the native S/390 volume type of the same name.

STI. See *self-timed interface*.

storage area network. A network that connects a company's heterogeneous storage resources.

storage complex. Multiple storage facilities.

storage device. A physical unit that provides a mechanism to store data on a given medium such that it can be subsequently retrieved. See also *disk drive module*.

storage facility. (1) A physical unit that consists of a storage server integrated with one or more storage devices to provide storage capability to a host computer. (2) A storage server and its attached storage devices.

storage server. A physical unit that manages attached storage devices and provides an interface between them and a host computer by providing the function of one or more logical subsystems. The storage server can provide functions that are not provided by the storage device. The storage server has one or more clusters.

striping. A technique that distributes data in bit, byte, multibyte, record, or block increments across multiple disk drives.

subchannel. A logical function of a channel subsystem associated with the management of a single device.

Subsystem Device Driver. See *IBM TotalStorage Enterprise Storage Server Subsystem Device Driver*.

subsystem identifier (SSID). A number that uniquely identifies a logical subsystem within a computer installation.

support catcher. See *catcher*.

support catcher telephone number. The telephone number that connects the support catcher server to the ESS to receive a trace or dump package. See also *support catcher* and *remote technical assistance information network*.

switched fabric. In the ESS, one of three fibre-channel connection topologies that the ESS supports. See also *arbitrated loop* and *point-to-point*.

symmetric multi-processor (SMP). An implementation of a multi-processor computer consisting of several identical processors configured in a way that any subset of the set of processors is capable of continuing the operation of the computer. The ESS contains four processors set up in SMP mode.

synchronous PPRC. A function of a storage server that maintains a consistent copy of a logical volume on the same storage server or on another storage server. All modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume. See also *Peer-to-Peer Remote Copy* and *PPRC Extended Distance*.

synchronous write. A write operation whose completion is indicated after the data has been stored on a storage device.

System/390. See *S/390*.

system adapter identification number (SAID). In the ESS, the unique identification number automatically assigned to each ESS host adapter for use by ESS Copy Services.

System Management Interface Tool (SMIT). An interface tool of the AIX operating system for installing, maintaining, configuring, and diagnosing tasks.

System Modification Program. A program used to install software and software changes on MVS systems.

T

TAP. See *Telocator Alphanumeric Protocol*.

target. A SCSI device that acts as a slave to an initiator and consists of a set of one or more logical units, each with an assigned logical unit number (LUN). The logical units on the target are typically I/O devices. A SCSI target is analogous to an S/390 control unit. A SCSI initiator is analogous to an S/390 channel. A SCSI logical unit is analogous to an S/390 device. See also *small computer system interface*.

TB. See *terabyte*.

TCP/IP. See *Transmission Control Protocol/Internet Protocol*.

Telocator Alphanumeric Protocol (TAP). An industry standard protocol for the input of paging requests.

terabyte (TB). (1) Nominally, 1 000 000 000 000 bytes, which is accurate when speaking of bandwidth and disk storage capacity. (2) For ESS cache memory, processor storage, real and virtual storage, a terabyte refers to 2^{40} or 1 099 511 627 776 bytes.

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

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

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

TotalStorage. See *IBM TotalStorage*.

TPF. See *transaction processing facility*.

track. A unit of storage on a CKD device that can be formatted to contain a number of data records. See also *home address*, *track-descriptor record*, and *data record*.

track-descriptor record (R0). A special record on a track that follows the home address. The control program uses it to maintain certain information about the track. The record has a count field with a key length of zero, a data length of 8, and a record number of 0. This record is sometimes referred to as R0.

transaction processing facility (TPF). A high-availability, high-performance IBM operating system, designed to support real-time, transaction-driven applications. The specialized architecture of TPF is intended to optimize system efficiency, reliability, and responsiveness for data communication and database processing. TPF provides real-time inquiry and updates to a large, centralized database, where message length is relatively short in both directions, and response time is generally less than three seconds. Formerly known as the Airline Control Program/Transaction Processing Facility (ACP/TPF).

Transmission Control Protocol (TCP). A communications protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

Transmission Control Protocol/Internet Protocol (TCP/IP). (1) A combination of data-transmission protocols that provide end-to-end connections between applications over interconnected networks of different types. (2) A suite of transport and application protocols that run over the Internet Protocol. (GC) See also *Internet Protocol* and *Transmission Control Protocol*.

transparency. See *software transparency*.

TSO. See *time sharing option*.

U

UFS. UNIX filing system.

Ultra-SCSI. An enhanced small computer system interface.

unconfigure. To delete the configuration.

unit address. In Enterprise Systems Architecture/390, the address associated with a device on a given control unit. On ESCON or FICON interfaces, the unit address is the same as the device address. On OEMI interfaces, the unit address specifies a control unit and device pair on the interface.

unprotected volume. An AS/400 term that indicates that the AS/400 host recognizes the volume as an unprotected device, even though the storage resides on a RAID array and is therefore fault tolerant by definition. The data in an unprotected volume can be mirrored. Also referred to as an *unprotected device*.

upper-layer protocol. The layer of the Internet Protocol (IP) that supports one or more logical protocols (for example, a SCSI-command protocol and an ESA/390 command protocol). Refer to ANSI X3.230-199x.

UTC. See *Coordinated Universal Time*.

utility device. The ESA/390 term for the device used with the Extended Remote Copy facility to access information that describes the modifications performed on the primary copy.

V

virtual machine facility. A virtual data processing machine that appears to the user to be for the exclusive use of that user, but whose functions are accomplished by sharing the resources of a shared data processing system. An alternate name for the VM/370 IBM operating system.

vital product data (VPD). Information that uniquely defines the system, hardware, software, and microcode elements of a processing system.

VM. The root name of several IBM operating systems, such as VM/370, VM/ESA, VM/CMS, and VM/SP. See also *virtual machine (VM) facility*.

volume. In Enterprise Systems Architecture/390, the information recorded on a single unit of recording medium. Indirectly, it can refer to the unit of recording medium itself. On a nonremovable-medium storage device, the term can also indirectly refer to the storage device associated with the volume. When multiple volumes are stored on a single storage medium transparently to the program, the volumes can be referred to as logical volumes.

VPD. See *vital product data*.

VSE/ESA. IBM operating system, the letters of which represent virtual storage extended/enterprise systems architecture.

W

Web Copy Services. See *ESS Copy Services*.

worldwide node name (WWNN). A unique 64-bit identifier for a host containing a fibre-channel port. See also *worldwide port name*.

worldwide port name (WWPN). A unique 64-bit identifier associated with a fibre-channel adapter port. It is assigned in an implementation- and protocol-independent manner.

write hit. A write operation in which the requested data is in the cache.

write penalty. The performance impact of a classical RAID 5 write operation.

WWPN. See *worldwide port name*.

X

XD. See *PPRC Extended Distance*.

XRC. See *Extended Remote Copy*.

xSeries. The product name of an IBM @server product that emphasizes industry-standard server scalability and self-managing server technologies. It is the successor to the Netfinity family of servers.

Z

z/Architecture. An IBM architecture for mainframe computers and peripherals. The IBM @server zSeries family of servers uses the z/Architecture architecture. It is the successor to the S/390 and 9672 family of servers. See also *Enterprise Systems Architecture/390*.

z/OS. An operating system for the IBM eServer product line that supports 64-bit real storage.

zSeries. (1) An IBM @server family of servers that emphasizes near-zero downtime. (2) IBM enterprise servers based on z/Architecture.

zSeries storage. Storage arrays and logical volumes that are defined in the ESS as connected to zSeries servers. See also *S/390 storage*.

Index

Special characters

/opt/IBMdpo/bin/showvpath command 152
/usr/lib/errstop command 25
> errclear 0 command 25
> errpt > file.save command 25

A

about this book xiii
accessing
  AIX
    add a data path volume to a volume group SMIT panel 57
    add a volume group with data path devices SMIT panel 57
    add paths to available data path devices SMIT panel 56
    back up a volume group with data path devices SMIT panel 58
    configure a defined data path device SMIT panel 56
    define and configure all data path devices SMIT panel 56
    display data path device adapter status SMIT panel 55
    Display Data Path Device Adapter Status SMIT panel 55
    display data path device configuration SMIT panel 55
    remake a volume group with data path devices SMIT panel 59
    remove a data path device SMIT panel 57
    Remove a Physical Volume from a Volume group SMIT panel 58
adapters
  configuring
    Linux 72
    Windows 2000 99
    Windows NT 83
  Emulex LP70000E 7
adding
  paths
    Windows 2000 host systems 103
    Windows NT 84
  storage for Windows NT host systems 88, 89
adding paths
  AIX
    to SDD devices volume group 20
  from AIX 4.3.2 volume group 20
addpaths
  utility programs, AIX 60
addpaths command 20, 52, 60
agreement for licensed internal code 168
AIX
  4.3.2
    32–bit applications 11, 12
    64–bit applications 11, 12
AIX *(continued)*
  4.3.3
    32–bit applications 11, 12
    64–bit applications 11, 12
  5.1.0
    32–bit applications 11, 12
    64–bit applications 11, 12
  accessing
    add a data path volume to a volume group SMIT panel 57
    add a volume group with data path devices SMIT panel 57
    add paths to available data path devices SMIT panel 56
    back up a volume group with data path devices SMIT panel 58
    configure a defined data path device SMIT panel 56
    define and configure all data path devices SMIT panel 56
    display data path device adapter status SMIT panel 55
    Display Data Path Device Adapter Status SMIT panel 55
    display data path device configuration SMIT panel 55
    remake a volume group with data path devices SMIT panel 59
    Remove a Physical Volume from a Volume group SMIT panel 58
    removing a data path device SMIT panel 57
  adding paths 20
  adding paths to SDD devices of a volume group 20
  applications
    32-bit 11, 12
    64-bit 11, 12
  backing-up files belonging to an SDD volume group 54
  changing the path-selection policy 19
  configuring
    volume group for failover protection 46
  configuring SDD 17
  error log messages 69
    new and modified messages by SDD for HACMP 70
  exporting
    volume group with SDD 49
  extending
    an existing SDD volume group 53
  importing
    volume group with SDD 48
  installation upgrade 21
  installing SDD 14
  migrating
    an existing non-SDD volume group to SDD vpath devices in concurrent mode 67
    non-SDD volume group to an ESS SDD multipath volume group in concurrent mode 65

© Copyright IBM Corp. 1999, 2002

AIX *(continued)*
recovering
from mixed volume groups 53
removing SDD from a host system 24
restoring files belonging to an SDD volume group 54
SDD utility programs 59
SDD-specific SMIT panels 54
the loss of a device path 50
unconfiguring SDD 17
upgrading 21
upgrading manually 23
upgrading SDD through a installation upgrade 23
verifying SDD 18
verifying SDD installation 15, 22
AIX 5.1.0
32–bit 12
64–bit 12
AIX fibre-channel requirements 8
AIX host system
disk driver 1
protocol stack 1
AIX levels
AIX 4.3.2
required PTFs 6
AIX 4.3.3
required maintenance level 6
AIX trace 69
article
Microsoft Knowledge Base Article Number Q293778
information about removing multipath access to your shared volume 93

B

backing-up AIX files belonging to an SDD volume group 54
BIOS, disabling 83, 99
block disk device interfaces (SDD) 112, 137
boot -r command 150
bootinfo -K command 12

C

Canadian compliance statement 167
cat /proc/modules command 74
cat /proc/scsi/scsi command 75
cat /proc/scsi/xxx/N command 75
cat /proc/sdd command 75
cd /mnt command 72
cd /opt/IBMsdd command 73
cd /opt/IBMsdd/bin command 73
cfgmgr
run *n* times where *n* represents the number of paths per SDD device. 52
run for each installed SCSI or fibre adapter 52
cfgmgr command 9, 21, 52
cfgvpath command 75
changing
path-selection policy for AIX 19
changing *(continued)*
SDD hardware configuration
HP host systems 118
Sun hosts 142
to the /dev directory
HP host systems 126
chdev command 51
chgrp command 133, 150
chkconfig - -level X sdd on command 76
chkconfig - -level *X* sdd off command 76
chkconfig - -list sdd command 76
chkvpenv command 74
chmod command 133, 150
chown command 133
class A compliance statement, Taiwan 168
command
/usr/lib/errstop 25
> errclear 0 25
> errpt > file.save 25
addpaths 20, 52, 60
bootinfo -K 12
cfgmgr 9, 52
running *n* times for n-path configurations 21, 52
running for each relevant SCSI or FCP adapter 21
chdev 51
datapath open device 162
datapath query adapter 154
datapath query adaptstats 155
datapath query device 46, 52, 156
datapath query devstats 158
datapath set adapter 160
datapath set device 161
datapath set device 0 path 0 offline 161
datapath set device N policy rr/fo/lb/df 20
dpovgfix 46, 60
dpovgfix vg-name 21, 51
extendvg 54
extendvg4vp 54, 61
hd2vp and vp2hd 60
hd2vp vg_name 24
HP host system
hd2vp 121
vgexport 122, 123
vgimport 123
vp2hd 122
installp 8
instfix -i | grep IY10201 8
instfix -i | grep IY10994 8
instfix -i | grep IY11245 8
instfix -i | grep IY13736 8
instfix -i | grep IYl7902 8
instfix -i | grep IYl8070 8
ls -al /unix 12
lscfg -vl fcsN 9
lsdev -Cc disk 9
lsdev -Cc disk | grep 2105 17
lslpp -l ibmSdd_432.rte 15, 16
lslpp -l ibmSdd_433.rte 15, 16
lslpp -l ibmSdd_510.rte 15, 16
lslpp -l ibmSdd_510nchacmp.rte 15, 16

command *(continued)*
lspv 21, 23, 49
lsvg -p vg-name 47
lsvgfs 23
lsvpcfg 21, 24, 45, 50, 60
mkdev -l vpath***N*** 18
mksysb restore command 50
mkvg 47
mkvg4vp 47, 61
odmget -q ″name = ioaccess″ CuAt 35
restvg 54
restvg4vp 54
rmdev 52, 53
rmdev -dl dpo -R 24
rmdev -dl fcs*N* -R 9
rmdev -l dpo -R 18
savevg 54
savevg4vp 54
shutdown -rF 9
smitty 24
smitty deinstall 8
smitty device 24
umount 23
using 153
varyoffvg 17, 23
varyonvg vg_name 24
commands
/opt/IBMdpo/bin/showvpath 152
boot -r 150
cat /proc/modules 74
cat /proc/scsi/scsi 75
cat /proc/scsi/xxx/N 75
cat /proc/sdd 75
cd /mnt 72
cd /opt/IBMsdd 73
cd /opt/IBMsdd/bin 73
cfgvpath 75
chgrp 133, 150
chkconfig - -level X sdd off 76
chkconfig - -level *X* sdd on 76
chkconfig - -list sdd 76
chkvpenv 74
chmod 133, 150
chown 133
datapath query
adapter 96, 109
device 76, 84, 101, 102, 109
datapath query device 115
datapath set adapter # offline 108
datapath set adapter offline 108
esvpcfg 79
insmod ./vpath.o 74
ls -l 73
lsvpcfg 76
lsvpsd 75
metadb -a <device> 151
metadb -d -f <device> 151
metadb -i 151
metainit 151
metainit d <metadevice number> -t <″vpathNs″ - master device> <″vpathNs″ - logging device> 152
commands *(continued)*
metastat 151, 152
newfs 152
orainst /m 132, 148
pkgrm IBMdpo 151
rmvpath xxx 75
rpm -e IBMsdd command 77
rpm -qi IBMsdd 73, 77
rpm -ql IBMsdd 73, 77
sdd restart 76
sdd start 76
sdd stop 76
showvpath 126, 148, 149, 150, 151
shutdown -i6 -y -g0 151
umount 152
umount /cdrom 141
unmod ./sdd-mod.o 77
commands datapath set adapter # offline 94
commands datapath set adapter offline 94
comments, how to send xxi
communications statement 166
compliance statement
German 167
radio frequency energy 166
Taiwan class A 168
concurrent download of licensed internal code
AIX 25
SDD 4
configuring
additional paths on a Windows NT host system 85
AIX
ESS 7
fibre-channel attached devices 8, 9
volume group for failover protection 46
clusters with SDD
Windows 2000 host system 108
Windows NT host system 95
ESS
HP host systems 112
Linux host systems 72
Sun host systems 136
Windows 2000 98
Windows NT 82
fibre-channel adapters
Linux host systems 72
Windows 2000 host system 99
Windows NT host system 82
SCSI adapters
Windows 2000 host systems 99
Windows NT 83
SDD
at system startup 76
Linux host systems 73, 74
Windows NT host system 84
SDD for AIX host 17
SDD on AIX 17
configuring a vpath device to the Available state 53
configuring all vpath devices to the Available state 53
conversion script
vp2hd 17

conversion scripts
  hd2vp 59
  vp2hd 21, 59
creating
  device node for the logical volume device in an HP host systems 125
  directory in /dev for the volume group in an HP host systems 125
  file system on the volume group in an HP host systems 126
  logical volume in an HP host systems 126
  new logical volumes in an HP host systems 125
  physical volume in an HP host systems 126
  volume group in an HP host systems 126
customizing
  Network File System file server 129
  Oracle 130, 146
  standard UNIX applications 124, 145

D

Data Path Optimizer (DPO) 83
database managers (DBMS) 137
datapath
  open device command 162
  query
    adapter command 96, 109
    device command 76, 84, 96, 109
  query adapter command 154
  query adaptstats command 155
  query device command 156
  query devstats command 158
  query set adapter command 160
  set adapter # offline command 94, 108
  set adapter offline command 94, 108
  set device command 161
datapath query device command 46, 52
datapath set device 0 path 0 offline command 161
datapath set device N policy rr/fo/lb/df command 20
definitions 171
determining
  AIX
    Emulex adapter firmware level 9
  major number of the logical volume device for an HP host systems 125
  size of the logical volume for an HP host systems 128
device driver 136
displaying
  AIX
    ESS vpath device configuration 45
  current version of SDD
    Windows 2000 101
  current version of the SDD
    Windows NT 93
documents, ordering xx
dpovgfix command 46, 60
dpovgfix vg-name command 21, 51
dynamic I/O load-balancing 3

E

electronic emission notices 166
Emulex adapter
  Emulex LP70000E 7
  firmware level 9
  upgrading firmware level to (sf320A9) 10
enhanced data availability 3
Enterprise Storage Server (ESS)
  library xv
  publications xv
error log messages
  AIX
    new and modified messages by SDD for HACMP 70
    VPATH_DEVICE_OFFLINE 69
    VPATH_DEVICE_ONLINE 69
    VPATH_PATH_OPEN 69
    VPATH_XBUF_NOMEM 69
error log messages for ibmSdd_433.rte installation package for SDD
  AIX
    VPATH_DEVICE_OPEN 70
    VPATH_FAIL_RELPRESERVE 70
    VPATH_OUT_SERVICE 70
    VPATH_RESV_CFLICT 70
error recovery
  HP-UX 11.0 host system 121
  HP-UX 11i host system 121
  SUN host system 144
error recovery and retry policy
  algorithms
    Windows 2000 host systems 106
    Windows NT host systems 93
error recovery and retry policy, for Windows 2000 host systems 106
error recovery and retry policy, for Windows NT host systems 93
error recovery policy
  algorithms
    Linux host systems 78
error recovery policy, for Linux host systems 78
ESS
  AIX
    displaying vpath device configuration 45
  configuring for HP 112
  configuring for Linux 72
  configuring for Sun 136
  configuring on
    Windows 2000 98
  configuring on Windows NT 82
ESS (Enterprise Storage Server)
  library xv
  publications xv
ESS devices (hdisks) 63
ESS LUNs 63
European Community Compliance statement 167
exporting a volume group with SDD, AIX 49
extending an existing SDD volume group, AIX 53
extendvg command 54
extendvg4vp command 54, 61

F

failover 3
failover protection, AIX
  creating a volume group from a single-path vpath device 50
  losing 49
  manually deleted devices and running the configuration manager 52
  providing load-balancing and failover protection 45
  side effect of running the disk change method 50
  the loss of a device path 50
  when it doesn't exist 45
Federal Communications Commission (FCC) statement 166
fibre-channel adapters
  configuring
    Linux host systems 72
  configuring for Windows 2000 99
  supported
    HP host systems 111
    Linux host systems 71
    Sun host systems 135
    Windows 2000 host systems 98
    Windows NT host systems 82
  supported on AIX host systems 7
fibre-channel device drivers
  configuring for AIX 8
  devices.common.IBM.fc 8
  devices.fcp.disk 8
  devices.pci.df1000f7 8
  installing for AIX 8
  supported on AIX host systems 7

G

German compliance statement 167
glossary 171

H

HACMP
  concurrent mode 30
  hd2vp conversion script 35
  node failover 43
  nonconcurrent mode 30
  persistent reserve 34
  recovering paths 43
  SDD installation package attributes 34
  software support for concurrent mode 31
  software support for nonconcurrent mode 32
  special requirements 35
  supported features 34
hardware configuration
  changing
    HP host systems 118
    Sun host systems 142
hardware requirements
  HP
    host systems 111
  Linux host systems 71
hardware requirements *(continued)*
  Sun host systems 135
hd2vp and vp2hd command 60
hd2vp command
  HP host system 121
hd2vp vg_name command 24
hdisk device
  chdev 50
  modify attributes 50
High Availability Cluster Multi-Processing (HACMP) 30
HP
  SCSI disk driver (sdisk) 112
HP host system
  disk driver 1
  protocol stack 1
HP host systems
  changing
    SDD hardware configuration 118
    to the /dev directory 126
  creating
    a file system on the volume group 126
    a logical volume 126
    a volume group 126
    device node for the logical volume device 125
    directory in /dev for the volume group 125
    new logical volumes 125
    physical volume 126
  determining
    major number of the logical volume 125
    size of the logical volume 128
  installing Oracle 131
  installing SDD 114
    converting an Oracle installation from sdisk 133
    on a Network File System file server 129
    on a system that already has Network File System file server 130
    on a system that already has Oracle 132
    using a file system 132
    using raw partitions 132
  mounting the logical volume 126
  recreating
    existing logical volume 127
    logical volume 128
    physical volume 126, 128
    volume group 128
  removing
    existing logical volume 127
    existing volume group 127
    logical volumes 127
  SDD 111
  setting the correct timeout value for the logical volume manager 129
  setting up Network File System for the first time 130
  setting up Oracle
    using a file system 131
    using raw partitions 131
  standard UNIX applications 124
  understanding how SDD works 112
  upgrading SDD 112, 117
  using applications with SDD 124

HP-UX
  disk device drivers 116, 124
  LJFS file system 130
  operating system 111
HP-UX 11.0
  32-bit 111, 113
  64-bit 111, 113
HP-UX 11i
  32-bit 111, 113
  64-bit 111, 113

I

ibm2105.rte 8
ibm2105.rte ESS package 6
ibmSdd_433.rte installation package
  for SDD 1.2.2.0
    removing 35
  for SDD 1.3.2.0. vpath devices
    unconfiguring 35
importing a volume group with SDD, AIX 48
Industry Canada Compliance statement 167
insmod ./sdd-mod.o command 74
installation package
  AIX 14
    ibmSdd_*nnn*.rte 14
    ibmSdd_432.rte 11, 12, 13, 16, 25, 30, 31, 65
    ibmSdd_433.rte 11, 12, 13, 16, 18, 25, 31, 34, 35, 48, 65, 70
    ibmSdd_510.rte 11, 12, 13, 16, 17, 25, 30, 31
    ibmSdd_510nchacmp.rte 11, 12, 13, 25, 30, 31
installation scenarios for SDD 1.3.3.3 99
installation upgrade
  AIX
    SDD 1.3.0.1 or later 21
installing
  additional paths on a Windows NT host system 85
  AIX
    fibre-channel device drivers 8
    planning 5
    SDD 14
  an Oracle installation from sdisk on an HP host systems 133
  converting an Oracle installation from sdisk on a Sun host system 149
  Oracle
    HP host systems 131
    Sun host systems 147
  SDD
    HP host systems 111, 114
    Linux host systems 71, 72
    Sun host systems 135, 138
    Windows 2000 host system 97
    Windows NT host system 81, 83
  SDD 1.3.3.3
    Windows 2000 host system 100
  SDD on a Network File System file server on a Sun host system 145
  SDD on a Network File System file server on an HP host systems 129
installing *(continued)*
  SDD on a system that already has Network File System file server 146
  SDD on a system that already has Network File System file server on an HP host systems 130
  SDD on a system that already has Oracle on a Sun host system 148
  SDD on a system that already has Oracle on an HP host systems 132
  SDD on a system that already has Solaris DiskSuite in place on a Sun host system 151
  Solstice DiskSuite for the first time on a Sun host system 150
  using a file system on a Sun host system 148
  using a file system on an HP host systems 132
  using raw partitions on a Sun host system 149
  using raw partitions on an HP host systems 132
  vpath on a system that already has UFS logging in place on a Sun host system 152
installp command 8
instfix -i | grep IY10201 command 8
instfix -i | grep IY10994 command 8
instfix -i | grep IY11245 command 8
instfix -i | grep IY13736 command 8
instfix -i | grep IYl7902 command 8
instfix -i | grep IYl8070 command 8

J

Japanese Voluntary Control Council for Interference (VCCI) statement 168

K

KB 158
Korean government Ministry of Communication (MOC) statement 168

L

library, ESS xv
licensed internal code
  agreement 168
Linux host system
  disk driver 1
  protocol stack 1
Linux host systems
  configuring
    ESS 72
    fibre-channel adapters 72
    SDD 73, 76
  error recovery policy 78
  installing SDD 71, 72
  maintaining vpath device configuration persistency 77
  partitioning SDD vpath devices 78
  preparing SDD installation 71
  removing SDD 77
  using
    multiple command for SDD configuration 74
    single command for SDD configuration 76

Linux host systems *(continued)*
using *(continued)*
standard UNIX applications 79
verifying SDD installation 73
load-balancing, AIX 45
loading SDD on Linux 73, 74
logical volume manager 137
losing failover protection, AIX 49
ls -al /unix command 12
ls -l command 73
lscfg -vl fcsN command 9
lsdev -Cc disk | grep 2105 command 17
lsdev -Cc disk command 9
lslpp -l *Sdd* command 15
lslpp -l ibmSdd_432.rte command 15, 16
lslpp -l ibmSdd_433.rte command 15, 16
lslpp -l ibmSdd_510.rte command 15, 16
lslpp -l ibmSdd_510nchacmp.rte command 15, 16
lspv command 21, 23, 49
lsvg -p vg-name command 47
lsvgfs command 23
lsvpcfg command 21, 24, 45, 50, 60, 76, 79
lsvpcfg utility programs, AIX 60
lsvpd command 75

M

maintaining vpath device configuration persistency, for Linux host systems 77
manuals, ordering xx
metadb -a <device> command 151
metadb -d -f <device> command 151
metadb -i command 151
metainit command 151
metainit d <metadevice number> -t <″vpathNs″ - master device> <″vpathNs″ - logging device> command 152
metastat command 151, 152
migrating
AIX
an existing non-SDD volume group to SDD vpath devices in concurrent mode 67
non-SDD volume group to an ESS SDD multipath volume group in concurrent mode 65
mirroring logical volumes 66
mkdev -l vpath*N* command 18
mksysb restore command 50
mkvg command 47
mkvg4vp command 47, 61
modifying multipath storage configuration to the ESS, Windows NT host system 88
mounting the logical volume, HP 126

N

newfs command 152
non-supported environments
AIX 6
HP 111
Linux 71
Windows NT 81
nonsupported environments
Sun 135
notices
electronic emission 166
European community 167
FCC statement 166
German 167
Industry Canada 167
Japanese 168
Korean 168
licensed internal code 168
notices statement 165
Taiwan 168

O

odmget -q ″name = ioaccess″ CuAt command 35
orainst /m command 132, 148
ordering publications xx

P

partitioning SDD vpath devices, for Linux host systems 78
path-failover protection system 3
path-selection policy
changing 20
default 20
failover only 19
load balancing 19
round robin 19
Persistent Reserve command set 34
pkgrm IBMdpo command 151
planning
AIX
Emulex adapter firmware level 9, 10
ESS 7
fibre-channel attached devices 8, 9
fibre-channel device drivers 8
preparing 7
AIX installation 5
ESS
HP host systems 112
Linux host systems 72
Sun host systems 136
Windows 2000 host system 98
Windows NT host system 82
fibre-channel adapters
Windows 2000 host system 99
Windows NT host system 82
hardware and software requirements on a Sun host system 135
hardware and software requirements on an HP host systems 111
hardware requirements, AIX
ESS 5
Fibre adapters and cables 5
Host system 5
SCSI adapters and cables 5
hardware requirements, Windows 2000
ESS 97

planning *(continued)*
hardware requirements, Windows NT
ESS 81
host system requirements, AIX 6
ESS 6
Fibre 7
SCSI 6
host system requirements, Windows 2000
ESS 97
host system requirements, Windows NT
ESS 81
installation of SDD
HP host systems 112
Sun host systems 137
preparing
Sun host systems 136
preparing for SDD installation on an HP host systems 112
SCSI adapters
Windows NT host systems 83
SDD
HP host systems 111
Linux host systems 71
Sun host systems 135
Windows 2000 host system 98
Windows NT host system 81
software requirements
Windows 2000 operating system 97
Windows NT operating system 81
software requirements, AIX
AIX operating system 6
ibm2105.rte ESS package 6
SCSI and fibre-channel device drivers 6
Windows 2000
ESS 98
planning for SDD 1.3.3.3 (or later) installation 99
post-installation of SDD
HP host systems 116
Sun host systems 140
preparing
AIX
SDD installation 7
configure on AIX 17
SDD
HP host systems 112
Linux host systems 71
Windows 2000 installation 98
Windows NT host system 82
SDD installation
Sun host systems 136
providing
AIX
failover protection 45
load-balancing 45
pSeries 690 with static LPARs configured 29
publications
ESS xv
ordering xx
related xvii
pvid 65
PVID 49

R

radio frequency energy compliance statement 166
raw
device interface (sd) 137
device interface (sdisk) 112
recovering from mixed volume groups 53
recovering from mixed volume groups, AIX 53
recovery procedures for HP 127, 129
recreating
existing logical volume
on a HP host systems 127
physical volume
on an HP host systems 126
the logical volume
on an HP host systems 128
the physical volume
on a HP host systems 128
the volume group
on an HP host systems 128
related publications xvii
removing
existing logical volume
on an HP host system 127
existing volume group
on an HP host systems 127
logical volumes
on an HP host systems 127
SDD
Linux host systems 77
Windows 2000 host system 103
Windows NT host system 92
SDD from an AIX host 24
SDD from an AIX host system 24
removing SDD in a 2-node cluster environment 103
requirements
ESS
Windows 2000 host system 97
Windows NT 81
hardware and software
Linux host systems 71
hardware and software on a Sun host system 135
hardware and software, HP 111
hardware, AIX
ESS 5
Fibre adapters and cables 5
Host system 5
SCSI adapters and cables 5
hardware, Windows 2000
ESS 97
hardware, Windows NT
ESS 81
host system, AIX 6
ESS 6
Fibre 7
SCSI 6
software
Windows 2000 operating system 97
Windows NT operating system 81
software, AIX
AIX operating system 6
ibm2105.rte ESS package 6

requirements *(continued)*
software, AIX *(continued)*
SCSI and fibre-channel device drivers 6
restoring
AIX
files belonging to an SDD volume group 54
restvg command 54
restvg4vp command 54
reviewing the existing SDD configuration information, Windows NT 84, 88
rmdev -dl dpo -R command 24
rmdev -dl fcsN -R command 9
rmdev -l dpo -R command 18
rmdev command 52, 53
rmvpath xxx command 75
rpm -e IBMsdd command 77
rpm -qi IBMsdd command 73, 77
rpm -ql IBMsdd command 73, 77

S

savevg command 54
savevg4vp command 54
SCSI adapter support
Windows 2000 host system 98
Windows NT host system 82
SCSI adapters
supported on AIX host systems 7
SCSI adapters support
HP host systems 111
Sun host systems 135
SCSI-3 Persistent Reserve command set 34
SDD
architecture 1
displaying the current version on Windows 2000 101
how it works on an HP host system 112
how it works on Sun 136
installation scenarios 112
installing
HP host system 111
Linux 71
Sun host systems 135
Windows 2000 host system 97, 100
Windows NT 81
installing on AIX 5
introducing 1
introduction 1
overview 1
post-installation of SDD
HP host systems 116
post-installation on Sun host systems 140
removing SDD on Windows NT 92
uninstalling
HP host systems 117
upgrading
HP host systems 112
Windows 2000 101
using applications
with SDD on HP Network File System file server 129
SDD *(continued)*
using applications *(continued)*
with SDD on HP standard UNIX applications 124
with SDD on HP, Oracle 130
with SDD on Linux standard UNIX applications 79
with SDD on Sun Network File System file Server 145
with SDD on Sun standard UNIX applications 145
with SDD on Sun, Oracle 146
verifying additional paths to SDD devices 86, 105
verifying configuration 18
Web site xiii
SDD configuration
checking 18
SDD devices
reconfiguring 21
sdd restart command 76
SDD server daemon
AIX host system 14, 26
HP host system 119
SUN host system 142
sdd start command 76
sdd stop command 76
SDD utility programs, AIX 59
SDD vpath devices 63
sddserver.rte
AIX host system 10
sddsrv
AIX host system 14, 26
HP host system 119
SUN host system 142
sddsrv for ESS Expert
AIX host system 10
setting up
correct timeout value for the logical volume manager on an HP host system 129
Network File System for the first time on a Sun host system 146
Network File System for the first time on an HP host system 130
Oracle using a file system
HP host system 131
Sun host system 147
Oracle using raw partitions
HP host system 131
Sun host system 147
UFS logging on a new system on a Sun host system 151
showvpath command 126, 148, 149, 150, 151
shutdown -i6 -y -g0 command 151
shutdown -rF command 9
smit
definition 14
SMIT
configuring
SDD for Windows NT host system 84
smitty
definition 14
smitty command 24

smitty deinstall command 8
smitty device command 24
software requirements
  for SDD on HP 111
  for SDD on Linux 71
  for SDD on Sun 135
Solaris
  host system
    upgrading the Subsystem Device Driver on 137
  operating system
    upgrading SDD 135
  sd devices 145
  UFS file system 146
statement
  of compliance
    Canada 167
    European 167
    Federal Communications Commission 166
    Japan 168
    Korean government Ministry of Communication (MOC) 168
    Taiwan 168
Subsystem device driver, see SDD. 150
Sun disk device drivers 137
Sun host system
  disk driver 1
  protocol stack 1
Sun host systems
  changing SDD hardware configuration 142
  installing
    Solstice DiskSuite for the first time 150
    vpath on a system that already has UFS logging in place 152
  installing Oracle 147
  installing SDD 138
    converting an Oracle installation from sdisk 149
    Network File System file server 145
    system that already has Network File System file server 146
    system that already has Oracle 148
    system that already has Solaris DiskSuite in place 151
    using a file system 148
    using raw partitions 149
  Oracle 146
  SDD 135
  SDD post-installation 140
  setting up
    Network File System for the first time 146
    UFS logging on a new system 151
  Solstice DiskSuite 150
  standard UNIX applications 145
  understanding how SDD works 136
  upgrading SDD 141
  using applications with SDD 145
Sun SCSI disk driver 136
support for Windows 2000 107
support for Windows NT 94
synchronizing logical volumes 66
System Management Interface Tool (SMIT) 14
  definition 14
System Management Interface Tool (SMIT) *(continued)*
  using for configuring 17
  using to access the Add a Data Path Volume to a Volume Group panel on AIX host 57
  using to access the Add a Volume Group with Data Path Devices panel on AIX host 57
  using to access the Back Up a Volume Group with Data Path Devices on AIX host 58
  using to access the Configure a Defined Data Path Device panel on AIX host 56
  using to access the Define and Configure All Data Path Devices panel on AIX host 56
  using to access the Display Data Path Device Adapter Status panel on AIX host 55
  using to access the Display Data Path Device Configuration panel on AIX host 55
  using to access the Remake a Volume Group with Data Path Devices on AIX host 59
  using to access the Remove a Data Path Device panel on AIX host 57
  using to access the Remove a Physical Volume from a Volume group panel on AIX host 58
  using to backup a volume group with SDD on AIX host 54, 58
  using to create a volume group with SDD on AIX host 46
  using to display the ESS vpath device configuration on AIX host 45
  using to export a volume group with SDD on AIX host 49
  using to extend an existing SDD volume group on AIX host 53
  using to import a volume group with SDD on AIX host 48
  using to remove SDD from AIX host 24
  using to restore a volume group with SDD on AIX host 54, 59
  using to unconfigure SDD devices on AIX host 17
  using to verify SDD configuration on AIX host 18

T

Taiwan class A compliance statement 168
trademarks 165

U

umount
  /cdrom command 141
  command 152
umount command 23
unconfiguring a SDD device to Defined state 53
unconfiguring all SDD devices to Defined state 53
unconfiguring SDD on AIX 17
understanding
  how SDD works for HP host systems 112
  how SDD works for Sun host systems 136
unmod ./sdd-mod.o command 77
upgrading
  AIX
    Emulex adapter firmware level 10

upgrading *(continued)*
AIX *(continued)*
manually 23
SDD through a installation upgrade 23
SDD
for AIX 4.3.2 21
for AIX 4.3.3 21
for AIX 5.1.0 21
HP host system 117
Sun host system 141
Windows 2000 host system 101
Windows NT host system 88
upgrading to SDD 1.3.3.3 (or later) in a 2-node cluster environment 102
using
HP applications with SDD 124
Linux
single command for SDD configuration 76
standard UNIX applications 79
Sun applications with SDD 145
using command 153
using ESS devices directly, AIX 63
using ESS devices through AIX LVM 64
using the datapath commands 153
using the trace function, AIX 69
utility programs, AIX
addpaths 60
dpovgfix 60
extendvg4vp 61
hd2vp and vp2hd 60
lsvpcfg 60
mkvg4vp 61
using ESS devices directly 63
using ESS devices through AIX LVM 64
using the trace function 69
utility programs, HP
hd2vp 121
vp2hd 122

V

varyoffvg command 17, 23
varyonvg vg_name command 24
verifying
additional paths are installed correctly
Windows 2000 host system 105
Windows NT host system 86
AIX
configuring SDD 18
SDD installation 15, 22
SDD installation
Linux host systems 73
vgexport command
HP host system 122, 123
vgimport command
HP host system 123
volume group
mixed
how to fix problem 51
mixed volume groups
dpovgfix vg-name 51
volume groups on AIX 47
vp2hd command
HP host system 122

W

Web site
AIX APARs, maintenance level fixes and microcode updates 6
information about removing multipath access to your shared volume
Multiple-Path Software May Cause Disk Signature to Change (Knowledge Base Article Number Q293778) 93
information on the fibre-channel adapters that can be used on your AIX host 7
information on the SCSI adapters that can attach to your AIX host 7
SDD xiii
Web sites
HP documentation 127, 129
information about
SCSI adapters that can attach to your Windows 2000 host system 98
SCSI adapters that can attach to your Windows NT host system 82
Windows 2000 host system
adding
paths to SDD devices 103
clustering special considerations 107
configuring
cluster with SDD 108
ESS 98
fibre-channel adapters 99
SCSI adapters 99
disk driver 1
displaying the current version of the SDD 101
installing SDD 1.3.3.3 100
path reclamation 107
protocol stack 1
removing SDD 103
SDD 97
support for clustering 107
upgrading SDD 101
verifying
additional paths to SDD devices 105
Windows 2000 host systems
error recovery and retry policy 106
Windows NT
adding
paths to SDD devices 84
Windows NT host system
adding
multipath storage configuration to the ESS 88
new storage to existing configuration 89
clustering special considerations 94
configuring
additional paths 85
clusters with SDD 95
SDD 84
disk driver 1

Windows NT host system *(continued)*
displaying the current version of the SDD 93
installing
additional paths 85
SDD 83
modifying multipath storage configuration to the ESS 88
path reclamation 94
protocol stack 1
removing SDD 92
reviewing existing SDD configuration information 84, 88
SDD 81
support for clustering 94
upgrading
SDD 88
verifying
additional paths to SDD devices 86
new storage is installed correctly 90
Windows NT host systems
error recovery and retry policy 93

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

IBM TotalStorage™ Enterprise Storage Server™:
Subsystem Device Driver User's Guide

Publication No. SC26-7478-00

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

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

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

Name _____ Address _____

Company or Organization _____

Phone No. _____

Readers' Comments — We'd Like to Hear from You
SC26-7478-00

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

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

Fold and Tape | **Please do not staple** | Fold and Tape

Cut or Fold
Along Line

Printed in U.S.A.

SC26-7478-00

Spine information:

IBM TotalStorage™ Enterprise Storage Server™:

SDD User's Guide