

IBM System Storage



Multipath Subsystem Device Driver User's Guide Errata

IBM System Storage



Multipath Subsystem Device Driver User's Guide Errata

Contents

Introduction vii

Who should use this book vii

Last update vii

Chapter 1. Using the SDD on an AIX host system 1

Supported SDD features 1

Verifying the hardware and software requirements 2

Hardware 2

Software 2

Unsupported environments 3

Host system requirements 3

Disk storage systems requirements 4

Virtualization products requirements 4

SCSI requirements for ESS 4

Fibre requirements 4

Preparing for the SDD installation 5

Configuring the disk storage system 5

Configuring the virtualization products 6

Installing the AIX fibre-channel device drivers 6

Uninstalling the AIX fibre-channel device drivers 7

Using the smitty deinstall command 7

Using the installp command 7

Installing the AIX SDD host attachment 8

Configuring fibre-channel-attached devices 8

Removing fibre-channel-attached devices 8

Verifying the adapter firmware level 9

Determining if the sddServer for Expert is installed 9

Understanding SDD support for IBM System p with static LPARs configured 10

Determining the installation package 10

Installation packages for 32-bit and 64-bit applications on AIX 4.3.3 (or later) host systems 10

Switching between 32-bit and 64-bit modes on AIX 5.1.0, AIX 5.2.0, and AIX 5.3.0 host systems 11

Installation of major files on your AIX host system 11

Determining the installation type 12

Installing and upgrading the SDD 13

Installing the SDD 13

Installing the SDD from CD-ROM 13

Installing SDD from downloaded code 14

Upgrading the SDD 14

Upgrading the SDD packages automatically without system restart 14

Preconditions for migration or upgrade 15

Customizing the SDD migration or upgrade 16

Procedures for automatic migration or upgrade 17

Error recovery for migration or upgrade 17

Upgrading SDD manually 18

Updating SDD packages by applying a program temporary fix 19

Committing or Rejecting a PTF Update 20

Upgrading AIX OS or host attachment and SDD packages 21

Case 1: In stand-alone host or with the HACMP services stopped 22

Case 2: In HACMP node-by-node migration with nonconcurrent resource group 23

Verifying the currently installed version of SDD for SDD 1.3.3.11 (or earlier) 25

Verifying the currently installed version of SDD for SDD 1.4.0.0 (or later) 26

Preparing to configure SDD 28

Maximum number of LUNs 28

ODM attributes for controlling the maximum number of LUNs in SDD version 1.6.0.7 or later on AIX 5.2 and later 29

Preparing your system to configure more than 600 supported storage devices or to handle a large amount of I/O after queue depth is disabled 30

Filesystem space 32

Controlling I/O flow to SDD devices with the SDD qdepth_enable attribute 32

Configuring SDD 33

Unconfiguring SDD 34

Verifying the SDD configuration 35

Dynamically adding paths to SDD vpath devices 36

Dynamically removing or replacing PCI adapters or paths 36

Dynamically removing a PCI adapter from SDD configuration 36

Dynamically replacing a PCI adapter in an SDD configuration 36

Dynamically removing a path of an SDD vpath device 38

Removing SDD from an AIX host system 38

Removing SDD Host Attachment from an AIX host system 40

SAN boot support 40

Manual exclusion of devices from the SDD configuration 40

Replacing manually excluded devices in the SDD configuration 41

SAN boot install procedures 41

SAN boot install procedure for AIX 5.1 41

SAN boot install procedure for AIX 5.2, AIX 5.3, and AIX 6.1 42

Understanding SDD support for High Availability Cluster Multiprocessing 43

SDD persistent reserve attributes 45

Preparation for importing volume groups under HACMP 46

HACMP RAID concurrent-mode volume groups and enhanced concurrent-capable volume groups 47

Creating HACMP RAID concurrent-mode volume groups	47	A side effect of running the disk change method	66
Importing HACMP RAID concurrent-mode volume groups	48	Manually deleting devices and running the configuration manager (cfgmgr)	67
Removing HACMP RAID concurrent-mode volume groups	50	Using LVM with SDD	68
Extending HACMP RAID concurrent-mode volume groups	51	Importing volume groups with SDD	68
Reducing HACMP RAID concurrent-mode volume groups	51	Exporting a volume group with SDD	69
Exporting HACMP RAID concurrent-mode volume groups	52	Recovering from mixed volume groups	69
Enhanced concurrent-capable volume groups	52	Extending an existing SDD volume group	70
Creating enhanced concurrent-capable volume groups	53	Backing up all files belonging to an SDD volume group	70
Importing enhanced concurrent-capable volume groups	53	Restoring all files belonging to an SDD volume group	70
Extending enhanced concurrent-capable volume groups	54	SDD-specific SMIT panels	71
Reducing enhanced concurrent-capable volume groups	55	Accessing the Display Data Path Device Configuration SMIT panel	72
Recovering paths that are lost during HACMP node fallover that is caused when a system locks up	55	Accessing the Display Data Path Device Status SMIT panel	72
Supporting enhanced concurrent mode in an HACMP environment	55	Accessing the Display Data Path Device Adapter Status SMIT panel	73
SDD server daemon	56	Accessing the Define and Configure All Data Path Devices SMIT panel	73
Verifying if the SDD server has started	56	Accessing the Add Paths to Available Data Path Devices SMIT panel	73
Starting the SDD server manually	57	Accessing the Configure a Defined Data Path Device SMIT panel	74
Changing to a different port number for the SDD server	57	Accessing the Remove a Data Path Device SMIT panel	74
Stopping the SDD server	57	Accessing the Add a Volume Group with Data Path Devices SMIT panel	74
Replacing the SDD server with a stand-alone version	57	Accessing the Add a Data Path Volume to a Volume Group SMIT panel	74
PTFs for APARs on AIX with Fibre Channel and the SDD server	58	Accessing the Remove a Physical Volume from a Volume Group SMIT panel	75
Dynamically changing the SDD path-selection policy algorithm	58	Accessing the Backup a Volume Group with Data Path Devices SMIT panel	75
datapath set device policy command	59	Accessing the Remake a Volume Group with Data Path Devices SMIT panel	75
Dynamically opening an invalid or close_dead path	59	SDD utility programs	76
Fibre-channel dynamic device tracking for AIX 5.20 TL1 (and later)	59	addpaths	76
Understanding SDD 1.3.2.9 (or later) support for single-path configuration for supported storage devices	61	hd2vp and vp2hd	76
Understanding the persistent reserve issue when migrating from SDD to non-SDD volume groups after a system reboot	61	dpovgfix	77
Managing secondary-system paging space	61	lsvpcfg	77
Listing paging spaces	62	mkvg4vp	78
Adding a paging space	62	extendvg4vp	78
Removing a paging space	62	querysn	79
Verifying load-balancing and failover protection	62	lquerypr	79
Displaying the supported storage device SDD vpath device configuration	62	sddgetdata	80
Configuring volume groups for failover protection	63	Persistent reserve command tool	80
Losing failover protection	65	Using supported storage devices directly	81
Losing a device path	66	Using supported storage devices through AIX LVM	82
Creating a volume group from single-path SDD vpath devices	66	Migrating a non-SDD volume group to a supported storage device SDD multipath volume group in concurrent mode	83
		Detailed instructions for migrating a non-SDD volume group to a supported storage device SDD multipath volume group in concurrent mode	85
		Using the trace function	86
		Chapter 2. Using SDDPCM on an AIX host system	89
		Supported SDDPCM features	91

Unsupported SDDPCM features	92	Dynamically adding and removing paths or adapters	107
Verifying the hardware and software requirements	92	Unconfiguring supported storage MPIO-capable devices	108
Hardware	92	Removing SDDPCM from an AIX host system	109
Software	92	SDDPCM support for HACMP with Enhanced Concurrent Mode volume groups	109
Unsupported environments	92	MPIO reserve policies	110
Host system requirements	93	No Reserve reservation policy	111
Supported storage device requirements	93	Exclusive Host Access single-path reservation policy	111
Fibre requirements	93	Persistent Reserve Exclusive Host Access reservation policy	111
Persistent Reserve Shared Host Access reservation policy	111	SDDPCM ODM attribute settings	111
Preparing for SDDPCM installation	93	SDDPCM ODM attribute default settings	111
Preparing for SDDPCM installation for supported storage devices	94	Changing device reserve policies	112
Determining the correct installation package	95	Changing the path selection algorithm	112
Determining if the SDD package is installed	95	Changing SDDPCM path healthcheck mode	113
Determining if the ibm2105.rte package is installed	95	Changing SDDPCM path healthcheck time interval	113
Determining if the devices.fcp.disk.ibm.rte package is installed.	95	AIX 5.2 TL07 (or later) or AIX 5.3 TL03 (or later) fibre-channel device driver features	114
Installing the AIX fibre-channel device drivers	95	Fast I/O failure of fibre-channel devices	114
Verifying and upgrading the fibre channel adapter firmware level	97	Fibre channel dynamic device tracking	115
Verifying the adapter firmware level	97	Multipath SAN boot support	116
Upgrading the adapter firmware level	97	Configuring supported storage system MPIO devices as the SAN boot device	116
Installing the AIX SDDPCM host attachment	97	Support system dump device with the supported storage system MPIO device	117
Installing and upgrading SDDPCM	98	Dynamically enabling and disabling paths or adapters	117
Installing SDDPCM from CD-ROM	98	Dynamically enabling or disabling a path	117
Creating and mounting the CD-ROM filesystem	98	Dynamically enabling or disabling an adapter	118
Using the System Management Interface Tool facility to install SDDPCM	100	Using the SDDPCM trace function	118
Unmounting the CD-ROM File System	100	SDDPCM server daemon	118
Installing SDDPCM from downloaded code	100	Verifying if the SDDPCM server has started	118
Installing SDDPCM with the AIX OS from an AIX NIM SPOT server to the client SAN boot disk or the internal boot disk	101	Starting the SDDPCM server manually	119
Updating SDDPCM	102	Stopping the SDDPCM server	119
Updating SDDPCM packages by installing a newer base package or a program temporary fix	102	SDDPCM utility programs	120
Committing or rejecting a program temporary fix update.	103	Persistent reserve command tools	120
Verifying the currently installed version of SDDPCM.	104	pcmquerypr	120
Maximum number of devices that SDDPCM supports	104	pcmgenprkey	122
Migrating SDDPCM	104	Using SDDPCM pcm _{path} commands	122
Migrating the supported storage SAN boot device or nonboot volume group from AIX default PCM to SDDPCM	104	pcm _{path} clear device count.	124
Migrating from SDDPCM to the AIX default PCM or to SDD	106	pcm _{path} disable ports	125
Migrating from SDD with SAN boot devices (on supported storage hdisks) to SDDPCM with multipath SAN boot devices	106	pcm _{path} enable ports	127
Migrating SDDPCM during an AIX OS upgrade with multipath SAN boot devices (on supported storage hdisks).	106	pcm _{path} open device path	129
Configuring and unconfiguring supported storage MPIO-capable devices	106	pcm _{path} query adapter	131
Configuring supported storage MPIO-capable devices	106	pcm _{path} query adaptstats	132
Verifying the SDDPCM configuration	107	pcm _{path} query device	133
		pcm _{path} query devstats	136
		pcm _{path} query essmap	138
		pcm _{path} query portmap.	139
		pcm _{path} query version	140
		pcm _{path} query wwpn	141
		pcm _{path} set adapter	142
		pcm _{path} set device algorithm	143
		pcm _{path} set device hc _{interval}	144
		pcm _{path} set device hc _{mode}	145

pcmpath set device path. 146
Summary of command syntax. 147

Index 149

Introduction

This guide provides errata information that pertains to release 1.6.3 of the *IBM® System Storage™ Multipath Subsystem Device Driver User's Guide*.

Specifically, this guide contains updates for the following two chapters that describe the SDD 1.7.0.x and SDDPCM 2.2.0.x drivers:

- Using the SDD on an AIX host system
- Using SDDPCM on an AIX host system

All of the changes in these chapters are indicated by a vertical bar (|).

Refer to the *IBM System Storage Multipath Subsystem Device Driver 1.6.3 User's Guide* for general information and if you are using the HP-UX, SUN Solaris, Linux, Windows, or Netware operating system and SDD 1.6.3.x.

Who should use this book

This errata is intended for those who are using AIX 6.1.0 (or later) and plan to install and configure IBM System Storage Multipath Subsystem Device Driver (SDD) 1.7.0 x or the IBM System Storage Multipath Subsystem Device Driver Path Control Module (SDDPCM) 2.2.0.x driver. The installer of SDD or SDDPCM should be:

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

Last update

This guide was last updated: November 15, 2007

Chapter 1. Using the SDD on an AIX host system

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

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

Starting from SDD 1.6.0.0, the SDD supports the coexistence of all supported storage devices.

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

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

Supported SDD features

The following SDD features are supported in this release:

- 32- and 64-bit kernels
- Support for ESS, DS8000, DS6000 and virtualization products
- Preferred node path-selection algorithm for DS6000 and virtualization products
- Changing the SDD path-selection algorithm dynamically. Five path-selection algorithms are supported:
 - Failover
 - Round robin
 - Round robin sequential
 - Load balancing
 - Load balancing sequential
- Dynamically adding paths to the SDD vpath devices
- Dynamically opening an invalid or close_dead path
- Dynamically removing or replacing PCI adapters or paths
- Fibre-channel dynamic device tracking
- SDD server daemon support
- Support for HACMP
- Support for secondary-system paging
- Support for load-balancing and failover protection for AIX applications and LVM
- SDD utility programs
- Support for SCSI-3 persistent reserve functions
- Support for AIX trace functions
- Support more than 512 SAN Volume Controller devices from multiple SAN Volume Controller clusters on an AIX host
- Storage I/O priority feature in DS6000 and DS8000, only with AIX53 TL04 or later and with 64-bit kernel
- Virtual I/O Server with AIX 5.3 or later

For more information about Virtual I/O Server, refer to the following Web site:
<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/home.html>

Verifying the hardware and software requirements

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

Hardware

The following hardware components are needed:

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

Software

The following software components are needed:

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

Packages for SDD 1.4.0.0 (and later) will be using new package names in order to comply with AIX packaging rules and allow for NIM installation. Table 1 shows the package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0 (or later).

Table 1. Package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0 (or later)

SDD 1.3.3.x	SDD 1.4.0.0 (or later)	Notes
ibmSdd_432.rte	N/A	Obsolete. This package has been merged with devices.sdd.43.rte.
ibmSdd_433.rte	devices.sdd.43.rte	N/A
ibmSdd_510.rte	N/A	Obsolete. This package has been merged with devices.sdd.51.rte.
ibmSdd_510nchacmp.rte	devices.sdd.51.rte	N/A
N/A	devices.sdd.52.rte	New package for AIX 5.2.0 (or later).
N/A	devices.sdd.53.rte	New package for AIX 5.3.0 (or later).

Table 1. Package-naming relationship between SDD 1.3.3.x and SDD 1.4.0.0 (or later) (continued)

SDD 1.3.3.x	SDD 1.4.0.0 (or later)	Notes
N/A	devices.sdd.61.rte	New package for AIX 6.1.0 (or later).

Notes:

1. SDD 1.4.0.0 (or later) no longer releases separate packages for concurrent and nonconcurrent High Availability Cluster Multiprocessing (HACMP). Both concurrent and nonconcurrent HACMP functions are now incorporated into one package for each AIX kernel level.
2. A persistent reserve issue arises when migrating from SDD to non-SDD volume groups after a reboot. This special case only occurs if the volume group was varied on prior to the reboot and auto varyon was not set when the volume group was created. See “Understanding the persistent reserve issue when migrating from SDD to non-SDD volume groups after a system reboot” on page 61 for more information.

Unsupported environments

The SDD does not support:

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

Host system requirements

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

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

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

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

www-03.ibm.com/servers/eserver/support/unixservers/aixfixes.html

Disk storage systems requirements

To successfully install the SDD:

Ensure that the disk storage system devices are configured as:

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

Virtualization products requirements

To successfully install the SDD, ensure that the SAN Volume Controller devices are configured as SAN Volume Controller Device (fibre-channel-attached device).

SCSI requirements for ESS

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

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

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

www.ibm.com/servers/eserver/support/unixservers/index.html

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.ibm.com/servers/eserver/support/unixservers/index.html

Notes:

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

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

www.ibm.com/servers/storage/support

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

- The AIX host system is an IBM RS/6000™ or IBM System p™ with AIX 4.3.3 (or later).
- The AIX host system has the fibre-channel device drivers installed along with all latest APARs.
- The bos.adt package is installed. The host system can be a single processor or a multiprocessor system, such as SMP.
- A fiber-optic cable connects each fibre-channel adapter to a disk storage system port.
- A fiber-optic cable connects each SAN Volume Controller fibre-channel adapter to a switch. The switch must also be configured correctly. Refer to the *IBM System Storage SAN Volume Controller: Software Installation and Configuration Guide* for information about the SAN Volume Controller.
- If you need the SDD I/O load-balancing and failover features, ensure that a minimum of two paths to a device are attached.

Preparing for the SDD installation

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

- Configuring the disk storage system
- Configuring the virtualization products
- Installing the AIX fibre-channel device drivers
- Installing the AIX SDD host attachment
- Configuring fibre-channel-attached devices
- Verifying the adapter firmware level
- Determining if the sddServer for Expert is installed
- Determining the installation package
- Determining the installation type

Note: The SDD allows the manual exclusion of supported devices from the SDD configuration. If you want to manually exclude supported devices (hdisks) from the SDD configuration, you must use the **querysn** command before configuring the SDD vpath devices. The **querysn** command reads the unique serial number of a device (hdisk) and saves the serial number in an exclude file. For detailed information about the **querysn** command, see “Manual exclusion of devices from the SDD configuration” on page 40.

Configuring the disk storage system

Before you install the SDD, you must configure:

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

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

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

Configuring the virtualization products

Before you install the SDD, you must configure:

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

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

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

Installing the AIX fibre-channel device drivers

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

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

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

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

devices.pci.df1000f9

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

devices.pci.df1000f7

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

devices.common.IBM.fc

The FCP protocol driver.

devices.fcp.disk

The FCP disk driver.

devices.pci.df1080f9

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

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

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

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

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

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

where *nnnnn* represents the APAR numbers.

If the APARs are listed, that means that they are installed. If they are installed, go to “Configuring fibre-channel-attached devices” on page 8. Otherwise, go to step 3.

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

Uninstalling the AIX fibre-channel device drivers

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

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

Using the smitty deinstall command

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

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

Using the installp command

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

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

Installing the AIX SDD host attachment

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

- Identify the hdisks as 2105, 2145, 2107, or a 1750 hdisks.
- Set default hdisk attributes, such as `queue_depth` and `timeout` values.
- Indicate to the AIX device driver configure method to configure the hdisk as a non-MPIO-capable device.

The following names are AIX SDD Host Attachment packages:

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

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

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

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

For installation instructions on attaching an AIX host, refer to the *Host Systems Attachment Guide* for your storage system.

Configuring fibre-channel-attached devices

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

- **cfgmgr** command

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

After the command prompt appears, use the **lsdev -Cc disk** command to check the Fibre Channel Protocol (FCP) disk configuration. If the FCP devices are configured correctly, they are in the *Available* state. If the FCP devices are configured correctly, go to “Verifying the adapter firmware level” on page 9 to determine if the proper firmware level is installed.

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

After the system restarts, use the **lsdev -Cc disk** command to check the Fibre Channel Protocol (FCP) disk configuration. If the FCP devices are configured correctly, they are in the *Available* state. If the FCP devices are configured correctly, go to “Verifying the adapter firmware level” on page 9 to determine if the proper firmware level is installed.

Removing fibre-channel-attached devices

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

```
rmdev -d1 fcs/ -R
```

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

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

Verifying the adapter firmware level

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

<http://www14.software.ibm.com/webapp/set2/firmware/gjsn>

Tip:

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

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

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

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

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

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

<http://www14.software.ibm.com/webapp/set2/firmware/gjsn>

Determining if the `sddServer` for Expert is installed

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

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

```
lspp -l sddServer.rte
```

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

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

For instructions on how to remove the stand-alone version of sddServer (for ESS Expert) from your AIX host system, see the IBM Subsystem Device Driver Server 1.0.0.0 (sddsrv) README for IBM TotalStorage Expert V2R1 at the following Web site:

www.ibm.com/servers/storage/support/software/swexpert/

For more information about the SDD server daemon, go to “SDD server daemon” on page 56.

Understanding SDD support for IBM System p with static LPARs configured

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

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

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

Determining the installation package

Before you install the SDD on your AIX host system (4.3.3 or later), you need to determine the installation package that is appropriate for your AIX OS level. See Table 2 to determine the correct installation package.

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

Table 2. SDD 1.4.0.0 (or later) installation packages for different AIX OS levels and the supported AIX kernel mode, application mode, and interface

SDD installation package names	AIX OS level	AIX kernel mode	Application mode	SDD interface
devices.sdd.43.rte	AIX 4.3.3 ¹	32-bit	32-bit, 64-bit	LVM, raw device

Table 2. SDD 1.4.0.0 (or later) installation packages for different AIX OS levels and the supported AIX kernel mode, application mode, and interface (continued)

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

¹devices.sdd.43.rte is supported only by the ESS and virtualization products.

²devices.sdd.61.rte supports only the 64-bit kernel.

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

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

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

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

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

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

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

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

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

Installation of major files on your AIX host system

The SDD installation package installs a number of major files on your AIX system. Table 3 lists the major files that are part of the SDD installation package.

Table 3. Major files included in the SDD installation package

File name	Description
-----------	-------------

Table 3. Major files included in the SDD installation package (continued)

defdpo	Define method of the SDD pseudo-parent data path optimizer (dpo).
cfgdpo	Configure method of the SDD pseudo-parent dpo.
define_vp	Define method of the SDD vpath devices.
addpaths	The command that dynamically adds more paths to SDD vpath devices while they are in <i>Available</i> state.
cfgvpath	Configure method of the SDD vpath devices.
chgvpath	Method to change vpath attributes.
cfallvpath	Fast-path configuration method to configure the SDD pseudo-parent dpo and all SDD vpath devices.
vpathdd	The SDD device driver.
hd2vp	The SDD script that converts an hdisk device volume group to an SDD vpath device volume group.
vp2hd	The SDD script that converts an SDD vpath device volume group to an hdisk device volume group.
datapath	The SDD driver console command tool.
lquerypr	The SDD driver persistent reserve command tool.
lsvpcfg	The SDD driver query configuration state command.
querysn	The SDD driver tool to query unique serial numbers of devices.
mkvg4vp	The command that creates an SDD volume group.
extendvg4vp	The command that extends the SDD vpath devices to an SDD volume group.
dpovgfix	The command that fixes an SDD volume group that has mixed vpath and hdisk physical volumes.
savevg4vp	The command that backs up all files belonging to a specified volume group with the SDD vpath devices.
restvg4vp	The command that restores all files belonging to a specified volume group with the SDD vpath devices.
sddsrv	The SDD server daemon for path reclamation and probe.
sample_sddsrv.conf	The sample SDD server configuration file.
lvmrecover	The SDD script that restores a system's SDD vpath devices and LVM configuration when a migration failure occurs.
sddfcmap	The SDD tool that collects information on ESS SCSI or disk storage systems fibre-channel devices through SCSI commands.
sddgetdata	The SDD data collection tool for problem analysis.

Determining the installation type

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

If there is no previous version of the SDD installed on the host system, see "Installing the SDD" on page 13 for instructions on installing and configuring the SDD. If there is a previous version of the SDD installed on the host system and you want to upgrade to one of the following packages:

- devices.sdd.43.rte
- devices.sdd.51.rte

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

See “Upgrading the SDD packages automatically without system restart” on page 14 for instructions on upgrading the SDD.

If SDD 1.4.0.0 (or later) is installed on the host system and you have an SDD PTF that you want to apply to the system, see “Updating SDD packages by applying a program temporary fix” on page 19 for instructions. A PTF file has a file extension of bff (for example, devices.sdd.43.rte.2.1.0.1.bff) and requires special consideration when being installed.

Installing and upgrading the SDD

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

Installing the SDD

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

Table 2 on page 10 lists the correct installation packages for different AIX OS level and the supported AIX kernel mode, application mode, and interface.

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

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

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

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

Installing the SDD from CD-ROM

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

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

8. Select **Software to Install** and press **F4**. The Software to Install panel is displayed.
9. Select the installation package that is appropriate for your environment.
10. Press **Enter**. The Install and Update from LATEST Available Software panel is displayed with the name of the software that you selected to install.
11. Check the default option settings to ensure that they are what you need.
12. Press **Enter** to install. SMIT responds with the following message:

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

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

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

Installing SDD from downloaded code

Use this procedure to install SDD from downloaded code:

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

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

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

6. Continue the installation by following the instructions beginning in step 3 on page 13.

Upgrading the SDD

Upgrading the SDD packages automatically without system restart

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

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

3. Package upgrade from version 1.4.0.0 to a later version.

If the SDD currently installed on your host system is listed in Table 4, you can use this automatic migration to upgrade the SDD. If the SDD currently installed on your host system is not listed in Table 4, you will need to upgrade the SDD manually

Table 4. List of previously installed installation packages that are supported with the installation upgrade

Installation package name
ibmSdd_432.rte
ibmSdd.rte.432
ibmSdd_433.rte
ibmSdd.rte.433
ibmSdd_510.rte
ibmSdd_510nchacmp.rte
devices.sdd.43.rte
devices.sdd.51.rte
devices.sdd.52.rte
devices.sdd.53.rte
devices.sdd.61.rte

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

Furthermore, the SDD now introduces two new environment variables that can be used in some configuration environments to customize the SDD migration and further reduce the time needed to migrate or upgrade the SDD. See “Customizing the SDD migration or upgrade” on page 16 for details.

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

Preconditions for migration or upgrade: The following are the preconditions for running the migration:

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

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

Customizing the SDD migration or upgrade: Starting from SDD 1.6.0.0, SDD offers two new environment variables, `SKIP_SDD_MIGRATION` and `SDDVG_NOT_RESERVED`, for you to customize the SDD migration or upgrade to maximize performance. You can set these two variables based on the configuration of your system. The following discussion explains the conditions and procedures for using these two environment variables.

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

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

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

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

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

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

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

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

Procedures for automatic migration or upgrade: To start the SDD migration or upgrade:

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

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

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

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

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

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

The following messages indicate that the postinstallation of SDD is successful:

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

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

Case 1: Smitty install failed.

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

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

If there is an error during the postinstallation (either the configuration of SDD devices has failed or LVM restoration has failed), the new SDD package is still successfully installed. Thus, the **Smitty install** exits with an OK prompt. However,

an error message indicating the error is printed, so you can identify and fix the problem. Then, run the shell script **lvmrecover** to configure SDD devices and automatically recover the original LVM configuration.

Upgrading SDD manually

The following section describes the procedure to migrate or upgrade SDD manually. See “Upgrading the SDD packages automatically without system restart” on page 14 for information about migrating or upgrading SDD automatically.

A manual migration or upgrade is required if you are upgrading from a previous version of the SDD package not listed in Table 4 on page 15. You must uninstall the existing SDD and then manually install the new version of SDD in these cases.

Perform the following steps to upgrade SDD:

1. Enter the **lspv** command to find out all the SDD volume groups.
2. Enter the **lsvgfs** command for each SDD volume group to find out which file systems are mounted. Enter the following command:
lsvgfs vg_name
3. Enter the **umount** command to unmount all file systems belonging to the SDD volume groups. Enter the following command:
umount filesystem_name
4. Enter the **varyoffvg** command to vary off the volume groups. Enter the following command:
varyoffvg vg_name
5. If you are upgrading to an SDD version earlier than 1.6.0.0; or if you are upgrading to SDD 1.6.0.0 or later and your host is in a HACMP environment with nonconcurrent volume groups that are varied-on on other host, that is, reserved by other host, run the **vp2hd volume_group_name** script to convert the volume group from the SDD vpath devices to supported storage hdisk devices. Otherwise, you skip this step.
6. Stop the SDD server by entering the following command:
stopsrc -s sddsrv
7. Remove all the SDD vpath devices. Enter the following command:
rmdev -dl dpo -R
8. Use the **smitty** command to uninstall the SDD. Enter **smitty deinstall** and press **Enter**. The uninstallation process begins. Complete the uninstallation process. See “Removing SDD from an AIX host system” on page 38 for the step-by-step procedure for uninstalling the SDD.
9. If you need to upgrade the AIX operating system, for example, from AIX 4.3 to AIX 5.1, you could perform the upgrade now. If required, reboot the system after the operating system upgrade.
10. Use the **smitty** command to install the newer version of the SDD from the compact disc. Enter **smitty install** and press **Enter**. The installation process begins. Go to “Installing and upgrading the SDD” on page 13 to complete the installation process.
11. Use the **smitty device** command to configure all the SDD vpath devices to the *Available* state. See “Configuring SDD” on page 33 for a step-by-step procedure for configuring devices.
12. Enter the **lsvpcfg** command to verify the SDD configuration. Enter the following command:
lsvpcfg

13. If you are upgrading to an SDD version earlier than 1.6.0.0, run the **hd2vp** *volume_group_name* script for each SDD volume group to convert the physical volumes from supported storage hdisk devices back to the SDD vpath devices. Enter the following command:
hd2vp *volume_group_name*
14. Enter the **varyonvg** command for each volume group that was previously varied offline. Enter the following command:
varyonvg *vg_name*
15. Enter the **lspv** command to verify that all physical volumes of the SDD volume groups are SDD vpath devices.
16. Enter the **mount** command to mount all file systems that were unmounted in step 3 on page 18. Enter the following command:
mount *filesystem-name*

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

Updating SDD packages by applying a program temporary fix

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

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

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

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

1. Log in as the root user.
2. Load the compact disc into the CD-ROM drive.
3. From your desktop window, enter **smitty install_update** and press **Enter** to go directly to the installation panels. The Install and Update Software menu is displayed.
4. Select **Install Software** and press **Enter**.
5. Press **F4** to display the INPUT Device/Directory for Software panel.
6. Select the compact disc drive that you are using for the installation (for example, */dev/cd0*) and press **Enter**.
7. Press **Enter** again. The Install Software panel is displayed.
8. Select **Software to Install** and press **F4**. The Software to Install panel is displayed.
9. Select the PTF package that you want to install.
10. Press **Enter**. The Install and Update from LATEST Available Software panel is displayed with the name of the software that you selected to install.

11. If you only want to apply the PTF, select **Commit software Updates?** and tab to change the entry to *no*. The default setting is to commit the PTF. If you specify *no* to **Commit Software Updates?**, be sure that you specify *yes* to **Save Replaced Files?**.
12. Check the other default option settings to ensure that they are what you need.
13. Press **Enter** to install. SMIT responds with the following message:

```

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

```

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

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

Committing or Rejecting a PTF Update: Before you reject a PTF update, you need to stop `sddsrvc` and remove all SDD devices. The following steps will guide you through this process. If you want to commit a package, you will not need to perform these steps. Follow these steps prior to rejecting a PTF update:

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

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

1. Log in as the root user.
2. From your desktop window, enter `smitty install` and press **Enter** to go directly to the installation panels. The Software Installation and Maintenance menu is displayed.
3. Select **Software Maintenance and Utilities** and press **Enter**.

4. Select **Commit Applied Software Updates** to commit the PTF or select **Reject Applied Software Updates** to reject the PTF.
5. Press **Enter**. The Commit Applied Software Updates panel is displayed or the Reject Applied Software Updates panel is displayed.
6. Select **Software name** and press **F4**. The software name panel is displayed.
7. Select the Software package that you want to commit or reject.
8. Check the default option settings to ensure that they are what you need.
9. Press **Enter**. SMIT responds with the following message:

```

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

```

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

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

After the procedure to reject a PTF update completes successfully:

1. Use the **smitty device** command to configure all the SDD vpath devices to the *Available* state. See “Configuring fibre-channel-attached devices” on page 8 for a step-by-step procedure for configuring devices.
2. Enter the **lsvpcfg** command to verify the SDD configuration. Enter the following command:
lsvpcfg
3. If you have downgraded to an SDD version earlier than 1.6.0.0, run the **hd2vp** script for each SDD volume group to convert the physical volumes from supported storage hdisk devices back to SDD vpath devices. Enter the following command:
hd2vp vg_name
4. Enter the **varyonvg** command for each volume group that was previously varied offline. Enter the following command:
varyonvg vg_name
5. Enter the **lspv** command to verify that all physical volumes of the SDD volume groups are SDD vpath devices.
6. Enter the **mount** command to mount all file systems that were unmounted in step 4. Enter the following command:
mount filesystem-name

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

7. Start the SDD Server. Enter the following command:
startsrc -s sddsrv

Upgrading AIX OS or host attachment and SDD packages

SDD provides different packages to match the AIX OS level. If an AIX system is going to be upgraded to a different OS level, you need to install the corresponding SDD package for that OS level. Automatic migration of an SDD package from an

earlier OS level to a later OS level after an OS upgrade is not supported. For example, automatic migration from devices.sdd.43.rte to devices.sdd.51.rte after an OS upgrade from AIX 4.3 to AIX 5.1, or automatic migration from devices.sdd.51.rte to devices.sdd.52.rte after an OS upgrade from AIX 5.1 to 5.2, is not supported.

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

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

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

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

If you want to upgrade SDD only, see “Upgrading the SDD packages automatically without system restart” on page 14 or “Upgrading SDD manually” on page 18.

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

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

1. Stop all activities related to SDD devices:
 - a. Stop applications running on SDD volume groups or file systems.
 - b. If your host is in an HACMP environment, stop the cluster services in an orderly manner.
 - c. If you are not upgrading AIX OS and you can reboot the system, go to Step 5 on page 23. Otherwise, continue to the next step.
 - d. If there is active paging space created with SDD devices, deactivate the paging space.
 - e. Use the **umount** command to unmount all file systems of SDD volume groups.
 - f. Vary off all SDD volume groups.
2. Remove SDD vpath devices using the **rmdev -dl dpo -R** command.
3. If *rootvg* is on local SCSI disks, follow these procedures. Otherwise, if *rootvg* is on a SAN boot disk, skip to step 4.
 - a. Remove hdisk devices using the following command:

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

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

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


4. If you are upgrading the OS, follow these procedures. Otherwise, if you are not upgrading the OS, skip to step 5.
 - a. Run **stopsrc -s sddsrv** to stop the sddsrv daemon.
 - b. Uninstall SDD.
 - c. Upgrade to the latest version of the host attachment, if required. The following are package names:
 - `ibm2105.rte` for 2105 devices
 - `devices.fcp.disk.ibm.rte` for 2145, 2107, and 1750 devices
 - d. If *rootvg* is on a SAN boot disk, restart the system.
 - e. Make sure no disk group is online except *rootvg*. Migrate the AIX OS level. The system automatically restarts at the end of migration.
 - f. Install SDD for the new AIX OS level.
 - g. Configure SDD vpath devices by running the **cfallvpath** command.
 - h. Continue to step 6.
5. If you are not upgrading the operating system, follow these steps.
 - a. Upgrade to the latest version of Host Attachment, if required. The following are Host Attachment Package names:
 - `ibm2105.rte` for 2105 devices
 - `devices.fcp.disk.ibm.rte` for 2145, 2107, and 1750 devices
 - b. After upgrading Host Attachment,
 - If *rootvg* is on a SAN boot disk, restart the system. Then skip the rest of the steps and follow the procedures in “Upgrading the SDD packages automatically without system restart” on page 14 to upgrade SDD, if required.
 - If *rootvg* is on local SCSI disks and you can restart the system, skip the rest of the steps and restart the system. Then follow the procedures in “Upgrading the SDD packages automatically without system restart” on page 14 to upgrade SDD, if required.
 - If *rootvg* is on local SCSI disks and you cannot restart the system, continue to the next step.
 - c. Upgrade to the latest version of SDD, if required.
 - d. Configure hdisks and SDD vpath devices by running the **cfgmgr** command.
6. If your new SDD version is earlier than 1.6.0.0, run the **hd2vp** command on all SDD volume groups. Otherwise, skip this step.
7. Resume all activities related to SDD devices:
 - a. If there was active paging space created with SDD devices, activate the paging space.
 - b. If your host was in an HACMP environment, start the cluster services.
 - c. Vary on all SDD volume groups.
 - d. Mount all file systems.
 - e. Start applications running on SDD volume groups or file systems.

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

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

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

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

1. Stop all activities related to SDD devices, if there are any.
 - a. Stop applications running on SDD volume groups or file systems.
 - b. If there is active paging space created with SDD devices, deactivate the paging space.
 - c. Use the **umount** command to unmount all file systems of SDD volume groups.
 - d. Vary off all SDD volume groups.

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

2. Remove the SDD vpath devices using the **rmdev -dl dpo -R** command.
3. Run **stopsrc -s sddsrv** to stop the sddsrv daemon.
4. Uninstall SDD.
5. If you do not want to upgrade OS, skip to step 10 on page 25. If you want to upgrade OS, go to the step 6.
6. If *rootvg* is on local SCSI disks, follow these procedures. Otherwise, if *rootvg* is on a SAN boot disk, skip to step 7.
 - a. Remove hdisk devices using the following command:

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

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

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

Reboot automatically starts at the end of migration.

8. If *rootvg* is on a local SCSI disk, follow these procedures. Otherwise, if *rootvg* is on a SAN boot disk, skip to Step 9 on page 25.

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

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

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

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

9. Upgrade to the latest version of Host Attachment, if required. The following are Host Attachment Package names:
  - `ibm2105.rte` for 2105 devices
  - `devices.fcp.disk.ibm.rte` for 2145, 2107, and 1750 devices
10. If `rootvg` is on a SAN boot disk, restart the system.
11. Install SDD. If you have migrated to a new AIX OS level, make sure you install the SDD for the new AIX OS level.
12. On the HACMP active node, run `varyonvg -bu volume group name` on all the SDD non-concurrent volume groups that are shared with the standby node.
13. On the HACMP standby node, perform the following steps:
  - a. Configure hdisks and the SDD vpath devices using one of the following options:
    - Run `cfgmgr -vl fcsX` for each fibre channel adapter and then run `cfallvpath`
    - Run `cfgmgr`
  - b. If your new SDD version is earlier than 1.6.0.0, run `hd2vp` on all SDD volume groups. Otherwise, skip this step.
  - c. Run `importvg -L volume group name physical volume name` to update any possible Object Data Manager (ODM) changes on a volume group.
14. On the HACMP active node, run `varyonvg volume group name` on all SDD non-concurrent volume groups that are shared with the standby node.

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

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

```
lspp -l '*Sdd*'
```

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

Alternatively, you can enter one of the following commands:

```
lspp -l ibmSdd_432.rte
```

```
lspp -l ibmSdd_433.rte
```

```
lspp -l ibmSdd_510.rte
```

```
lspp -l ibmSdd_510nchacmp.rte
```

```
lspp -l ibmSdd.rte.432
```

...

...

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

| Fileset                                   | Level   | State     | Description                                      |
|-------------------------------------------|---------|-----------|--------------------------------------------------|
| Path: /usr/lib/objrepos<br>ibmSdd_432.rte | 1.3.3.9 | COMMITTED | IBM SDD<br>AIX V432 V433 for concurrent<br>HACMP |
| Path: /etc/objrepos<br>ibmSdd_432.rte     | 1.3.3.9 | COMMITTED | IBM SDD<br>AIX V432 V433 for concurrent<br>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<br>ibmSdd_433.rte | 1.3.3.9 | COMMITTED | IBM SDD<br>AIX V433 for nonconcurrent<br>HACMP |
| Path: /etc/objrepos<br>ibmSdd_433.rte     | 1.3.3.9 | COMMITTED | IBM SDD<br>AIX V433 for nonconcurrent<br>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<br>ibmSdd_510.rte | 1.3.3.9 | COMMITTED | IBM SDD<br>AIX V510 for concurrent HACMP |
| Path: /etc/objrepos<br>ibmSdd_510.rte     | 1.3.3.9 | COMMITTED | IBM SDD<br>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<br>ibmSdd_510nchacmp.rte | 1.3.3.11 | COMMITTED | IBM SDD<br>AIX V510 for nonconcurrent<br>HACMP |
| Path: /etc/objrepos<br>ibmSdd_510nchacmp.rte     | 1.3.3.11 | COMMITTED | IBM SDD<br>AIX V510 for nonconcurrent<br>HACMP |

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

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

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

Alternatively, you can enter one of the following commands:

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

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

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

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

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

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

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

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

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

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

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

---

## Preparing to configure SDD

Before you configure SDD, ensure that:

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

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

Perform the following steps:

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

## Maximum number of LUNs

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

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

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

Table 5 on page 29 provides a summary of the maximum number of LUNs and the maximum number of paths allowed when running on a host systems with different operating system levels.

Table 5. Maximum LUNs allowed for different AIX OS levels

| OS level  | SDD supported storage devices                                                      |
|-----------|------------------------------------------------------------------------------------|
| AIX 4.3 * | 600 LUNs (maximum 32 paths)                                                        |
| AIX 5.1   | 600 LUNs (maximum 32 paths)                                                        |
| AIX 5.2   | 1200 LUNs (maximum 32 paths; see Table 6 for recommended maximum number of paths.) |
| AIX 5.3   | 1200 LUNs (maximum 32 paths; see Table 6 for recommended maximum number of paths.) |
| AIX 6.1   | 1200 LUNs (maximum 32 paths; see Table 6 for recommended maximum number of paths.) |

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

You can have a maximum of 32 paths per SDD vpath device regardless of the number of LUNs configured. However, configuring more paths than is needed for failover protection might consume too many system resources and degrade system performance. You should use the minimum number of paths necessary to achieve sufficient redundancy in the SAN environment. The recommended number of paths is 2 - 4.

To avoid exceeding the maximum number of paths per SDD vpath device on AIX 5.2 or later, follow the recommendations in Table 6.

Table 6. Recommended maximum paths supported for different number of LUNs on AIX 5.2 or later

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

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

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

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

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

- 2105\_max\_luns
- 2145\_max\_luns
- 2062\_max\_luns
- Enterpr\_maxlun
- Virtual\_maxlun



See Table 5 on page 29 for information about the total number of LUNs that you can configure.

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

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

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

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

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

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

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

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

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

### `lg_term_dma`

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

### `num_cmd_elems`

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

|                        |      |
|------------------------|------|
| <b>LP9000 adapters</b> | 2048 |
|------------------------|------|



|                         |      |
|-------------------------|------|
| <b>LP10000 adapters</b> | 2048 |
| <b>LP7000 adapters</b>  | 1024 |

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

#### **max\_xfer\_size**

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

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

#### **fc\_err\_recov**

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

#### **Notes:**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

```
lsattr -El vpath0
active_hdisk hdisk66/13AB2ZA1020/fscsi3 Active hdisk False
active_hdisk hdisk2/13AB2ZA1020/fscsi2 Active hdisk False
active_hdisk hdisk34/13AB2ZA1020/fscsi2 Active hdisk False
active_hdisk hdisk98/13AB2ZA1020/fscsi3 Active hdisk False
policy df Scheduling Policy True
pvid 0005f9fdcda4417d0000000000000000 Physical volume identifier False
qdepth_enable yes Queue Depth Control True
serial_number 13AB2ZA1020 LUN serial number False
unique_id yes Device Unique Identification False
```

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

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

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

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

See “Preparing your system to configure more than 600 supported storage devices or to handle a large amount of I/O after queue depth is disabled” on page 30 to determine whether the system has sufficient resources for disabling queue depth logic.

---

## Configuring SDD

Perform the following steps to configure SDD using SMIT:

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

1. Enter **smitty device** from your desktop window. The Devices menu is displayed.
2. Select **Data Path Device** and press **Enter**. The Data Path Device panel is displayed.
3. Select **Define and Configure All Data Path Devices** and press **Enter**. The configuration process begins.
4. Check the SDD configuration state. See “Displaying the supported storage device SDD vpath device configuration” on page 62.
5. Use the **varyonvg** command to vary on all deactivated supported storage device volume groups.
6. Mount the file systems for all volume groups that were previously unmounted.

## Unconfiguring SDD

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

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

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

- If you unconfigure without deleting the device information, the device remains in the *Defined* state. You can use either SMIT or the **mkdev -l vpathN** command to return the device to the *Available* state.
- If you unconfigure and delete the device information from the ODM database, that device is removed from the system. To reconfigure it, follow the procedure described in “Configuring SDD” on page 33.

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

1. Enter **smitty device** from your desktop window. The Devices menu is displayed.
2. Select **Devices** and press **Enter**.
3. Select **Data Path Device** and press **Enter**. The Data Path Device panel is displayed.
4. Select **Remove a Data Path Device** and press **Enter**. A list of all SDD devices and their states (either *Defined* or *Available*) is displayed.
5. Select the device that you want to unconfigure. Select whether or not you want to delete the device information from the ODM database.
6. Press **Enter**. The device is unconfigured to the state that you selected.

7. To unconfigure more SDD devices, you have to repeat steps 4 - 6 for each SDD device.

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

## 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:

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

If any device is listed as *Defined*, the configuration was not successful. Check the configuration procedure again. See “Configuring SDD” on page 33 for the procedure.

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

You should see an output similar to this:

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

The output shows:

- The name of each SDD vpath device (for example, vpath13)
- The *Defined* or *Available* state of an SDD vpath device
- Whether or not the SDD vpath device is defined to AIX as a physical volume (indicated by the pv flag)
- The name of the volume group the device belongs to (for example, vpathvg)
- The unit serial number of the disk storage system LUN (for example, 02FFA067) or the unit serial number of the virtualization product LUN (for example, 60056768018A0210B00000000000006B)

- The names of the AIX disk devices making up the SDD vpath device and their configuration and physical volume state

## Dynamically adding paths to SDD vpath devices

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

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

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

You can determine how many more hdisks can be added to the existing SDD vpath devices by using the command **datapath query device** to find out how many hdisks are already configured for each vpath. Ensure that the number of hdisks from the existing configuration is below the recommended maximum paths allowed according to Table 6 on page 29.

### Notes:

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

## Dynamically removing or replacing PCI adapters or paths

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

- “Dynamically removing a PCI adapter from SDD configuration”
- “Dynamically replacing a PCI adapter in an SDD configuration”
- “Dynamically removing a path of an SDD vpath device” on page 38

### Dynamically removing a PCI adapter from SDD configuration

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

### Dynamically replacing a PCI adapter in an SDD configuration

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



After you physically replace and configure a new adapter, the adapter and its associated paths can be added to SDD with the **addpaths** command. See “datapath remove adapter” in Chapter 13 of the *IBM System Storage Multipath Subsystem Device Driver User’s Guide* for more information about the **datapath remove adapter n** command.

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

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

```
-----+-----
Active Adapters :4
-----+-----
```

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

```
-----+-----
```

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

```
-----+-----
Success: remove adapter 0
-----+-----
```

```
Active Adapters :3
-----+-----
```

| Adpt# | Adapter Name | State  | Mode   | Select | Errors | Paths | Active |
|-------|--------------|--------|--------|--------|--------|-------|--------|
| 1     | fscsi1       | NORMAL | ACTIVE | 65916  | 3      | 10    | 10     |
| 2     | fscsi2       | NORMAL | ACTIVE | 76197  | 28     | 10    | 10     |
| 3     | fscsi3       | NORMAL | ACTIVE | 4997   | 39     | 10    | 10     |

```
-----+-----
```

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

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

```
-----+-----
Active Adapters :4
-----+-----
```

| Adpt# | Adapter Name | State  | Mode   | Select | Errors | Paths | Active |
|-------|--------------|--------|--------|--------|--------|-------|--------|
| 0     | fscsi0       | NORMAL | ACTIVE | 11     | 0      | 10    | 10     |
| 1     | fscsi1       | NORMAL | ACTIVE | 196667 | 6      | 10    | 10     |

```
-----+-----
```

|   |        |        |        |        |    |    |    |
|---|--------|--------|--------|--------|----|----|----|
| 2 | fscsi2 | NORMAL | ACTIVE | 208697 | 36 | 10 | 10 |
| 3 | fscsi3 | NORMAL | ACTIVE | 95188  | 47 | 10 | 10 |

### Dynamically removing a path of an SDD vpath device

To dynamically remove a particular path from an SDD vpath device, use the **datapath remove device *m* path *n*** command. This command permanently removes the logical path from the SDD device. See “datapath remove device path” in Chapter 13 of the *IBM System Storage Multipath Subsystem Device Driver User’s Guide* for more information about the **datapath remove device *m* path *n*** command.

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

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

```

=====
DEV#: 0 DEVICE NAME: vpath0 TYPE: 2105E20 POLICY: Optimized
SERIAL: 20112028
=====
Path# Adapter/Hard Disk State Mode Select Errors

0 fscsi1/hdisk18 OPEN NORMAL 557 0
1 fscsi1/hdisk26 OPEN NORMAL 568 30
2 fscsi0/hdisk34 OPEN NORMAL 566 0
3 fscsi0/hdisk42 OPEN NORMAL 545 0
=====

```

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

```

=====
Success: device 0 path 1 removed
=====
DEV#: 0 DEVICE NAME: vpath0 TYPE: 2105E20 POLICY: Optimized
SERIAL: 20112028
=====
Path# Adapter/Hard Disk State Mode Select Errors

0 fscsi1/hdisk18 OPEN NORMAL 567 0
1 fscsi0/hdisk34 OPEN NORMAL 596 0
2 fscsi0/hdisk42 OPEN NORMAL 589 0
=====

```

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

3. To reclaim the removed path, see “Dynamically adding paths to SDD vpath devices” on page 36.

```

=====
DEV#: 0 DEVICE NAME: vpath0 TYPE: 2105E20 POLICY: Optimized
SERIAL: 20112028
=====
Path# Adapter/Hard Disk State Mode Select Errors

0 fscsi1/hdisk18 OPEN NORMAL 588 0
1 fscsi0/hdisk34 OPEN NORMAL 656 0
2 fscsi0/hdisk42 OPEN NORMAL 599 0
3 fscsi1/hdisk26 OPEN NORMAL 9 0
=====

```

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

## Removing SDD from an AIX host system

The SDD server (sddsrv) is an integrated component of SDD 1.3.2.9 (or later). The SDD server daemon is automatically started after SDD is installed. You must stop the SDD server if it is running in the background *before removing SDD*. Go to “Verifying if the SDD server has started” on page 56 and “Stopping the SDD



server” on page 57 for more instructions. See Chapter 12 of the *IBM System Storage Multipath Subsystem Device Driver User’s Guide* for more details about the SDD server daemon.

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

1. Stop all activities that are related to SDD devices, if there are any.
  - a. Stop applications that are running on SDD volume groups or file systems.
  - b. If there is active paging space that is created with SDD devices, deactivate the paging space.
  - c. Issue the **umount** command to unmount all file systems of SDD volume groups.
  - d. Vary off all SDD volume groups.
2. Remove the SDD vpath devices using the **rmdev -dl dpo -R** command.
3. SDD has a server daemon running in the background. Therefore, before uninstalling SDD, issue the **stopsrc -s sddsrv** command to stop the sddsrv daemon.
4. Perform the following steps to uninstall the SDD package.
  - a. Enter **smitty deinstall** from your desktop window to go directly to the Remove Installed Software panel.
  - b. Enter one of the following installation package names in the **SOFTWARE name** field:
    - devices.sdd.43.rte**
    - devices.sdd.51.rte**
    - devices.sdd.52.rte**
    - devices.sdd.53.rte**
    - devices.sdd.61.rte**Then press **Enter**.

**Note:** See “Verifying the currently installed version of SDD for SDD 1.3.3.11 (or earlier)” on page 25 or “Verifying the currently installed version of SDD for SDD 1.4.0.0 (or later)” on page 26 to verify your currently installed installation package or version of SDD. You can also press **F4** in the Software name field to list the currently installed installation package and do a search (/) on SDD.

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

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

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

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

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

---

## Removing SDD Host Attachment from an AIX host system

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

1. Remove hdisk devices using the following command:

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

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

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

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

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

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

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

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

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

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

---

## SAN boot support

These sections provide information about SAN boot support.

### Manual exclusion of devices from the SDD configuration

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

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

SDD 1.3.3.9 (or later) allows you to manually exclude devices from the SDD configuration. See “querysn” on page 79 for the syntax of the **querysn** command. The maximum number of devices that can be excluded is 100.

You can also exclude any virtualization products devices from the SDD configuration with the **querysn** command.

**Notes:**

1. You should not use the **querysn** command on the same logical device multiple times.
2. Fibre-channel boot capability is available for disk storage system only.

## Replacing manually excluded devices in the SDD configuration

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

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

1. If the excluded devices belong to an active volume group and file systems of that volume group are mounted, you need to perform one of the following actions:
  - a. Use the **umount** command to unmount all the file systems of the volume group and vary off the volume group.
  - b. Or, unmount all the file systems of the volume group and use the **reducevg** command to reduce that device from the volume group.
2. Use a text editor such as vi to open the `/etc/vpexclude` file and delete the line containing the device name from the file.
3. Issue **cfallvpath** configure methods to configure these new devices.
4. Issue **lsvpcfg** to verify that these devices are configured as SDD vpath devices.

## SAN boot install procedures

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

**Note:** Because of current AIX boot device limitation, no more than 5 physical volumes (hdisks) should be selected for a boot device.

### SAN boot install procedure for AIX 5.1

Use this procedure for SAN boot install for AIX 5.1:

1. Configure disk storage system devices to the AIX system; there should be only a single path per LUN. In other words, the AIX system should see only one hdisk configured per LUN.
2. Install the base OS on the selected disk storage system single-path devices.
3. Upgrade the base OS to latest technology level.
4. Install both SDD Host Attachment and SDD.
5. Reboot the AIX system.
6. Ensure that SDD vpath devices are configured correctly. Disk storage system devices (hdisks) should be configured as IBM 2105, IBM 2107, or IBM 1750 devices. Run the **datapath query device** command to verify that SDD vpath devices are configured with single-path and that the SDD vpath device policy is Single Path.
7. Issue **bosboot -aDd /dev/ipldevice**.
8. Connect additional host adapters and additional storage adapter to the fabric and configure multipath (multiple hdisks) per disk storage system LUN.
9. Reboot the AIX system.
10. Verify that SDD vpath devices are configured correctly with multipath per LUN.
11. Issue **bosboot -aDd /dev/ipldevice**.
12. Reboot the AIX system.

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

### **SAN boot install procedure for AIX 5.2, AIX 5.3, and AIX 6.1**

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

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

- c. Verify that all the logical device names of the hdisks (except rootvg) are configured in a continuous sequence.

---

## Understanding SDD support for High Availability Cluster Multiprocessing

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

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

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

For SDD versions earlier than 1.3.1.3, refer to the corresponding User's Guide for HACMP support information. Table 7 provides a list of installation packages and the SDD versions (earlier than 1.4.0.0) and the HACMP modes that support them.

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

*Table 7. Recommended SDD installation packages and supported HACMP modes for SDD versions earlier than SDD 1.4.0.0*

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

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

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

- | <http://www.ibm.com/systems/storage/disk/ess/>
- | <http://www.ibm.com/systems/storage/disk/ds6000/>
- | <http://www.ibm.com/systems/storage/disk/ds8000/>
- | <http://www.ibm.com/systems/storage/software/virtualization/svc/>

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

| Table 8 and Table 9 summarize the software requirements to support HACMP v4.5:

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

*Table 8. Software support for HACMP 4.5 on AIX 4.3.3 (32-bit only), 5.1.0 (32-bit and 64-bit), 5.2.0 (32-bit and 64-bit)*

| SDD version and release level                                                                                                                                                                                           | HACMP 4.5 + APARs                                                                                                                                                             |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| devices.sdd.43.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)                                                                                                                                           | Not applicable                                                                                                                                                                |
| devices.sdd.51.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)                                                                                                                                           | <ul style="list-style-type: none"> <li>• IY36938</li> <li>• IY36933</li> <li>• IY35735</li> <li>• IY36951</li> </ul>                                                          |
| devices.sdd.52.rte installation package for SDD 1.4.0.0 (or later) (SCSI/FCP)                                                                                                                                           | <ul style="list-style-type: none"> <li>• IY36938</li> <li>• IY36933</li> <li>• IY36782</li> <li>• IY37744</li> <li>• IY37746</li> <li>• IY35810</li> <li>• IY36951</li> </ul> |
| <p><b>Note:</b> For up-to-date APAR information for HACMP, go to the following Web site:<br/> <a href="https://techsupport.services.ibm.com/server/aix.fdc">https://techsupport.services.ibm.com/server/aix.fdc</a></p> |                                                                                                                                                                               |

*Table 9. Software support for HACMP 4.5 on AIX 5.1.0 (32-bit and 64-bit kernel)*

| SDD version and release level                                                          | HACMP 4.5 + APARs                                                                                                    |
|----------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| ibmSdd_510nchacmp.rte installation package for SDD 1.3.1.3 (SCSI/FCP)                  | <ul style="list-style-type: none"> <li>• IY36938</li> <li>• IY36933</li> <li>• IY35735</li> <li>• IY36951</li> </ul> |
| ibmSdd_510nchacmp.rte installation package for SDD 1.3.2.9 (to SDD 1.3.3.x) (SCSI/FCP) | <ul style="list-style-type: none"> <li>• IY36938</li> <li>• IY36933</li> <li>• IY35735</li> <li>• IY36951</li> </ul> |

Table 9. Software support for HACMP 4.5 on AIX 5.1.0 (32-bit and 64-bit kernel) (continued)

| SDD version and release level                                                                                                                                                                                           | HACMP 4.5 + APARs |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| <p><b>Note:</b> For up-to-date APAR information for HACMP, go to the following Web site:<br/> <a href="https://techsupport.services.ibm.com/server/aix.fdc">https://techsupport.services.ibm.com/server/aix.fdc</a></p> |                   |

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

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

For HACMP up-to-date ARAP information, go to the following Web site:

<http://www14.software.ibm.com/webapp/set2/sas/f/hacmp/download/aix53.html>

## SDD persistent reserve attributes

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

The SDD installation packages have a new attribute under the pseudo-parent (dpo) that reflects whether or not the supported storage device supports the Persistent Reserve Command set. The attribute name is *persistent\_resv*. If SDD detects that the storage device supports Persistent Reserve, the *persistent\_resv* attribute is created in the CuAt ODM and the value is set to *yes*; otherwise this attribute exists only in the PdAt ODM and the value is set to *no* (default). You can use the following command to check the *persistent\_resv* attribute after the SDD device configuration is complete:

```
lsattr -El dpo
```

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

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

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

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

The output should look similar to this:

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



## Preparation for importing volume groups under HACMP

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

Before SDD 1.6.0.0, SDD does not automatically create the pvid attribute in the ODM database for each SDD vpath device. The AIX disk driver automatically creates the pvid attribute in the ODM database, if a pvid exists on the physical device. Because SDD versions earlier than 1.6.0.0 do not automatically create the pvid attribute in the ODM database for each SDD vpath device, the first time that you import a new SDD volume group to a new cluster node, you must import the volume group using hdisks as physical volumes. Next, run the hd2vp conversion script (see "SDD utility programs" on page 76) to convert the volume group's physical volumes from supported storage device hdisks to SDD vpath devices. This conversion step not only creates pvid attributes for all SDD vpath devices that belong to that imported volume group, it also deletes the pvid attributes for the underlying hdisks for these SDD vpath devices. Later on, you can import and vary on the volume group directly from the SDD vpath devices. These special requirements apply to both concurrent and nonconcurrent volume groups.

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

There are four scenarios:

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

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

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

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

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

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

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

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



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

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

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

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

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

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

**Note:** See “Importing volume groups with SDD” on page 68 for a detailed procedure for importing a volume group with the SDD devices.

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

This section provides information about HACMP RAID concurrent-mode volume groups and enhanced concurrent-capable volume groups. This section also provides instructions on the following procedures for both HACMP RAID concurrent-mode volume groups and enhanced concurrent-capable volume groups.

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

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

This section describes the advantage of enhanced concurrent-capable volume groups in an HACMP environment. It also describes the different ways of creating two kinds of concurrent-capable volume groups. While there are different ways to create and vary on concurrent-capable volume groups, the instructions to export a volume group are always the same. See “Exporting HACMP RAID concurrent-mode volume groups” on page 52.

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

### Creating HACMP RAID concurrent-mode volume groups

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

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

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

```

Add a Volume Group with Data Path Devices

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

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

```

## Importing HACMP RAID concurrent-mode volume groups

When importing the volume group to other nodes in the cluster, you need to vary off the volume group on the node after it is created. You can import the volume group from either the SDD `vpath` device or the `hdisk` device, depending on the `pvid` condition on the node to which the volume group is to be imported. Follow this procedure to import a volume group with SDD `vpath` device.

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

```

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

```

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

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

```

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

```

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

```

Import a Volume Group

```

Type or select values in the entry fields.

Press Enter AFTER making all desired changes.

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

```

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

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

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

```

Import a Volume Group

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

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

```

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

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

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

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

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

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

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

```

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

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

```

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

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

```

Import a Volume Group

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

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

```

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

```

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

```

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

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

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

## Removing HACMP RAID concurrent-mode volume groups

### Notes:

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

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

1. Ensure that the volume group is varied on.
2. Enter **smitty vg** at the command prompt.

3. Select **Remove a Volume Group** from the displayed menu.

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

```

Remove a Volume Group

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

[Entry Fields]
VOLUME GROUP name [con_vg]

```

## Extending HACMP RAID concurrent-mode volume groups

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

1. Vary off the HACMP RAID concurrent-mode volume group to be extended on all nodes.
2. Enter **smitty datapath\_extendvg** at the command prompt of one of the nodes.
3. A screen similar to the following example is displayed. Enter the information appropriate for your environment. The following example shows how to extend an HACMP RAID concurrent-mode volume group using the *con\_vg* on an SDD *vpath2*.

```

Add a Datapath Physical Volume to a Volume Group

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

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

```

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

## Reducing HACMP RAID concurrent-mode volume groups

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

1. Vary off the HACMP RAID concurrent-mode volume group to be reduced on all nodes.
2. Enter **smitty vg** at the command prompt.
3. Select **Set Characteristics of a Volume Group** from the displayed menu.

4. Select **Remove a Physical Volume from a Volume Group** from the displayed menu.
5. A screen similar to the following example is displayed. Enter the information appropriate for your environment. The following example shows how to reduce an HACMP RAID concurrent-mode volume group using the *con\_vg* on an SDD *vpath1*. Assume that *con\_vg* originally has *vpath0* and *vpath1* as its physical volumes.

```

Remove a Physical Volume from a Volume Group

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

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

```

6. Vary off the volume group after reducing it on the current node.
7. For all the nodes sharing *con\_vg*, do the following:
  - a. Enter **exportvg con\_vg** at the command prompt.
  - b. Enter **smitty importvg** at the command prompt.
  - c. A screen similar to the following is displayed. Enter the information appropriate for your environment.

```

Import a Volume Group

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

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

```

- d. Verify that *con\_vg* has the *vpath* reduced by using the **lspv** command.

## Exporting HACMP RAID concurrent-mode volume groups

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

### Notes:

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

## Enhanced concurrent-capable volume groups

With the AIX51 TL02 and HACMP 4.4.1.4 environments, enhanced concurrent mode is supported with both 32-bit and 64-bit kernels. The advantage of this mode is that after you create an enhanced concurrent-capable volume group on multiple nodes, the changes made to the logical volume or volume group structures on one

node (for example, extending or reducing a volume group), are propagated to all other nodes. Also, the Logical Volume Manager (LVM) configuration files are updated on all nodes.

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

For more detailed information on enhanced concurrent-capable volume groups, see “Supporting enhanced concurrent mode in an HACMP environment” on page 55.

### Creating enhanced concurrent-capable volume groups:

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

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

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

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

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

```

Add a Volume Group with Data Path Devices

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

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

### Importing enhanced concurrent-capable volume groups:

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

Before importing enhanced concurrent-capable volume groups on SDD `vpath` devices, issue the `lspv` command to make sure there is `pvid` on the SDD `vpath`

device. If pvid is not displayed, import the volume group on one of the hdisks that belongs to the SDD vpath device. Enter **hd2vp** to convert the volume group to SDD vpath devices.

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

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

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

```

Import a Volume Group

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

[Entry Fields]
VOLUME GROUP name [con_vg]
PHYSICAL VOLUME names [vpath3]
Volume Group MAJOR NUMBER [45]
Make this VOLUME GROUP concurrent-capable? yes
Make default varyon of VOLUME GROUP concurrent? no

```

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

**Extending enhanced concurrent-capable volume groups:**

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

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

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

```

Add a Datapath Physical Volume to a Volume Group

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

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

```

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



## Reducing enhanced concurrent-capable volume groups:

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

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

```

Remove a Physical Volume from a Volume Group

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

[Entry Fields]
VOLUME GROUP name [con_vg]
PHYSICAL VOLUME names [vpath2]
FORCE deallocation of all partitions yes

```

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

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

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

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

If your SDD vpath devices have lost one or more underlying paths that belong to an active volume group, you can use either the Add Paths to Available Data Path Devices SMIT panel or run the **addpaths** command from the AIX command line to recover the lost paths. Go to “Dynamically adding paths to SDD vpath devices” on page 36 for more information about the **addpaths** command.

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

## Supporting enhanced concurrent mode in an HACMP environment

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

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

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

For more detailed information and instructions on creating, removing, reducing, importing, and exporting enhanced concurrent-capable volume groups, see “Enhanced concurrent-capable volume groups” on page 52.

---

## SDD server daemon

The SDD server (sddsrv) is an integrated component of SDD 1.3.2.9 (or later). This component consists of a UNIX application daemon that is installed in addition to the SDD device driver. See Chapter 12 of the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for more information about sddsrv.

**Attention:** Running sddsrv will activate several AIX Fibre Channel Protocol or adapter driver problems on AIX 4.3.3 and 5.1.0. One of the problems in the AIX Fibre Channel Protocol driver is that internal resources can be leaked. You will experience this as a performance degradation that grows worse over time. Performance can be restored by unconfiguring and reconfiguring the fibre-channel adapter or by rebooting the system. AIX users with Fibre Channel Protocol support and the SDD server daemon installed should apply the PTFs listed in “PTFs for APARs on AIX with Fibre Channel and the SDD server” on page 58.

## Verifying if the SDD server has started

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

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

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

where *NNN* is the process ID number.

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

If the SDD server has *not* started, the status will be *Inoperative*. Go to “Starting the SDD server manually” on page 57 to proceed.

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

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

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

## Starting the SDD server manually

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

Go to “Verifying if the SDD server has started” on page 56 to verify that the SDD server started successfully.

## Changing to a different port number for the SDD server

See “Changing the sddsrv or pcmsrv TCP/IP port number” in Chapter 12 of the *IBM System Storage Multipath Subsystem Device Driver User’s Guide*.

## Stopping the SDD server

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

If sddsrv is stopped, the feature that automatically recovers failed paths (DEAD or CLOSE\_DEAD path) is disabled. During the concurrent storage bay quiesce/resume process, you must manually recover the adapter or paths after the quiesce/resume is completed on one bay, and before the quiesce/resume starts on the next bay. Without doing so, the application might fail. See “datapath set device path” in Chapter 13 of the *IBM System Storage Multipath Subsystem Device Driver User’s Guide* for more information.

## Replacing the SDD server with a stand-alone version

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

1. Run **stopsrc -s sddsrv** to stop the current version of sddsrv.
2. Verify that the SDD server has stopped. See “Verifying if the SDD server has started” on page 56 and the status should be *inoperative*.
3. Comment out following line:

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

in the system /etc/inittab table.

4. Add following line:

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

to the system /etc/inittab table.

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

Starting sddsrv with the s0 flag does not provide path health check or path reclamation functions. You should manually recover paths by using the **datapath** command. See “datapath set device path” in Chapter 13 of the *IBM System Storage Multipath Subsystem Device Driver User’s Guide* for more information.

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

If you have fibre-channel support and the SDD server daemon running, PTFs for the APARs shown in Table 10 must be applied to your AIX servers in order to avoid a performance degradation.

Table 10. PTFs for APARs on AIX with fibre-channel support and the SDD server daemon running

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

If you experience a degradation in performance, you should disable sddsrv until the PTFs for these APARs can be installed. After the PTFs for these APARs are installed, you should re-enable sddsrv. If you are running IBM TotalStorage Expert, refer to “Replacing the SDD server with a stand-alone version” on page 57. Otherwise, refer to “Stopping the SDD server” on page 57.

---

## Dynamically changing the SDD path-selection policy algorithm

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

### failover only (fo)

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

### load balancing (lb)

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

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

### load balancing sequential (lbs)

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

### round robin (rr)

The path to use for each I/O operation is chosen at random from paths

that were not used for the last I/O operation. If a device has only two paths, SDD alternates between the two.

#### **round robin sequential (rrs)**

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

#### **default (df)**

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

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

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

```
lsattr -El vpath0
active_hdisk hdisk66/13AB2ZA1020/fscsi3 Active hdisk False
active_hdisk hdisk2/13AB2ZA1020/fscsi2 Active hdisk False
active_hdisk hdisk34/13AB2ZA1020/fscsi2 Active hdisk False
active_hdisk hdisk98/13AB2ZA1020/fscsi3 Active hdisk False
policy df Scheduling Policy True
pvid 0005f9fdcda4417d000000000000000 Physical volume identifier False
qdepth_enable no Queue Depth Control True
serial_number 13AB2ZA1020 LUN serial number False
unique_id yes Device Unique Identification False
```

## **datapath set device policy command**

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

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

See “datapath set device policy” in Chapter 13 of the *IBM System Storage Multipath Subsystem Device Driver User’s Guide* for more information about the **datapath set device policy** command.

---

## **Dynamically opening an invalid or close\_dead path**

With SDD 1.3.2.9 (or later), you can issue the **datapath open path** command to dynamically open a path that is in an INVALID or CLOSE\_DEAD state if the SDD vpath device it belongs to is in the OPEN state. You can use this command even when the I/O is actively running. See “datapath open device path” in Chapter 13 of the *IBM System Storage Multipath Subsystem Device Driver User’s Guide* for more information.

---

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

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

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

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

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

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

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

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

1. Issue the **rmdev -l fscsiX -R** for all adapters on your system to change all the child devices of fscsiX on your system to the defined state.
2. Issue the **chdev -l fscsiX -a dyntrk=yes** command for all adapters on your system.
3. Run **cfgmgr** to reconfigure all devices back to the available state.

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

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

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

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

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

---

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

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

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

**Tip:** It is also possible to change from single-path to multipath configuration by using the **addpaths** command. For more information about the **addpaths** command, go to “Dynamically adding paths to SDD vpath devices” on page 36.

---

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

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

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

There are two ways to avoid this issue.

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

---

## Managing secondary-system paging space

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

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

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

**Note:** AIX does not recommend moving the primary paging space out of rootvg. Doing so might mean that *no* paging space is available during the system



startup, which can result in poor startup performance. *Do not* redefine your primary paging space using SDD vpath devices.

## Listing paging spaces

You can list paging spaces by entering:

```
lsps -a
```

## Adding a paging space

You can add a paging space by entering:

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

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

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

## Removing a paging space

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

For example, to remove paging space PS01, enter:

```
rmpps PS01
```

---

## Verifying load-balancing and failover protection

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

## Displaying the supported storage device SDD vpath device configuration

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

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

1. Enter **smitty device** from your desktop window. The Devices panel is displayed.
2. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.



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

You will see an output similar to the following example:

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

The following information is displayed:

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

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

- The name of the volume group to which the device belongs, such as vpathvg.
- The unit serial number of the supported storage device LUN, such as 019FA067.
- The names of the AIX disk devices that comprise the SDD vpath devices, their configuration states, and the physical volume states. See “lsvpcfg” on page 77 for information about the **lsvpcfg** command.

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

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

The sample output shows that device vpath10 has two paths and both are operational. See “datapath query device” in Chapter 13 of the *IBM System Storage Multipath Subsystem Device Driver User’s Guide* for more information about the **datapath query device** command.

## Configuring volume groups for failover protection

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

It is possible to create a volume group that has only a single path (see 63) and then add paths later by reconfiguring the supported storage device. (See “Dynamically adding paths to SDD vpath devices” on page 36 for information about adding paths to an SDD device.) However, an SDD volume group does not have failover protection if any of its physical volumes has only a single path.

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

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

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

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

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

After you create the volume group, AIX creates the SDD vpath device as a physical volume (pv). In the output shown on page 63, vpath9 through vpath13 are included in a volume group and they become physical volumes. To list all the physical volumes known to AIX, use the **lspv** command. Any SDD vpath devices that were created into physical volumes are included in the output similar to the following output:

|         |                  |         |
|---------|------------------|---------|
| hdisk0  | 0001926922c706b2 | rootvg  |
| hdisk1  | none             | None    |
| ...     |                  |         |
| hdisk10 | none             | None    |
| hdisk11 | 00000000e7f5c88a | None    |
| ...     |                  |         |
| hdisk48 | none             | None    |
| hdisk49 | 00000000e7f5c88a | None    |
| vpath0  | 00019269aa5bc858 | None    |
| vpath1  | none             | None    |
| vpath2  | none             | None    |
| vpath3  | none             | None    |
| vpath4  | none             | None    |
| vpath5  | none             | None    |
| vpath6  | none             | None    |
| vpath7  | none             | None    |
| vpath8  | none             | None    |
| vpath9  | 00019269aa5bbadd | vpathvg |
| vpath10 | 00019269aa5bc4dc | vpathvg |
| vpath11 | 00019269aa5bc670 | vpathvg |
| vpath12 | 000192697f9fd2d3 | vpathvg |
| vpath13 | 000192697f9fde04 | vpathvg |

To display the devices that comprise a volume group, enter the `lsvg -p vg-name` command. For example, the `lsvg -p vpathvg` command might produce the following output:

| PV_NAME | PV STATE | TOTAL PPs | FREE PPs | FREE DISTRIBUTION  |
|---------|----------|-----------|----------|--------------------|
| vpath9  | active   | 29        | 4        | 00..00..00..00..04 |
| vpath10 | active   | 29        | 4        | 00..00..00..00..04 |
| vpath11 | active   | 29        | 4        | 00..00..00..00..04 |
| vpath12 | active   | 29        | 4        | 00..00..00..00..04 |
| vpath13 | active   | 29        | 28       | 06..05..05..06..06 |

The example output indicates that the `vpathvg` volume group uses physical volumes vpath9 through vpath13.

---

## Losing failover protection

AIX can create volume groups only from or SDD vpath devices that are physical volumes. If a volume group is created using a device that is not a physical volume, AIX makes it a physical volume as part of the procedure of creating the volume group. A physical volume has a physical volume identifier (pvid) written on its sector 0 and also has a pvid attribute attached to the device attributes in the CuAt ODM. The `lspv` command lists all the physical volumes known to AIX. Here is a sample output from this command:

|         |                  |         |
|---------|------------------|---------|
| hdisk0  | 0001926922c706b2 | rootvg  |
| hdisk1  | none             | None    |
| ...     |                  |         |
| hdisk10 | none             | None    |
| hdisk11 | 00000000e7f5c88a | None    |
| ...     |                  |         |
| hdisk48 | none             | None    |
| hdisk49 | 00000000e7f5c88a | None    |
| vpath0  | 00019269aa5bc858 | None    |
| vpath1  | none             | None    |
| vpath2  | none             | None    |
| vpath3  | none             | None    |
| vpath4  | none             | None    |
| vpath5  | none             | None    |
| vpath6  | none             | None    |
| vpath7  | none             | None    |
| vpath8  | none             | None    |
| vpath9  | 00019269aa5bbadd | vpathvg |
| vpath10 | 00019269aa5bc4dc | vpathvg |
| vpath11 | 00019269aa5bc670 | vpathvg |
| vpath12 | 000192697f9fd2d3 | vpathvg |
| vpath13 | 000192697f9fde04 | vpathvg |

In some cases, access to data is not lost, but failover protection might not be present. Failover protection can be lost in several ways:

- Losing a device path
- Creating a volume group from single-path SDD vpath devices
- A side effect of running the disk change method
- Running the `mksysb restore` command
- Manually deleting devices and running the configuration manager (cfmgr)

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

## Losing a device path

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

## Creating a volume group from single-path SDD vpath devices

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

## A side effect of running the disk change method

It is possible to modify attributes for an hdisk device by running the **chdev** command. The **chdev** command invokes the hdisk configuration method to make the requested change. In addition, the hdisk configuration method sets the pvid attribute for an hdisk if it determines that the hdisk has a pvid written on sector 0 of the LUN. This causes the SDD vpath device and one or more of its hdisks to have the same pvid attribute in the ODM. If the volume group containing the SDD vpath device is activated, the LVM uses the first device it finds in the ODM with the required pvid to activate the volume group.

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

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

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

```
chdev -l hdisk46 -a pv=yes
```

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

```

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

```

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

```

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

```

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

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

## Manually deleting devices and running the configuration manager (cfgmgr)

In AIX 4.3 and 5.1, if there is more than one adapter and the storage device has been used as a physical volume of a volume group, running the **cfgmgr** command once might not configure all the hdisks. This means that if you manually delete SDD vpath devices and its hdisks (paths) and run configuration manager, you might lose the failover protection.

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

```

rmdev -d1 vpath3
rmdev -d1 hdisk4
rmdev -d1 hdisk27
cfgmgr

```

The **datapath query device** command might show now that only one path (either hdisk4 or hdisk27) is configured for vpath3.

To restore failover protection (that is, configure multiple paths for vpath3), perform the following steps:

1. Enter **cfgmgr** once for each installed SCSI or fibre-channel adapter, or enter **cfgmgr** *n* times, where *n* represents the number of paths per SDD device.  
**Tip:** Running **cfgmgr** *n* times for *n*-path vpath configurations is not always required. It is only necessary to run **cfgmgr** *n* times for an *n*-path configuration if the supported storage device has been used as a physical volume of a volume group. This is because the AIX disk driver might configure only one set of hdisks from one adapter if pvid is present on a device.
2. Run **addpaths** to dynamically add the paths discovered by **cfgmgr** to SDD vpath devices.

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

This command opens a new path (or multiple paths) automatically if the SDD vpath device is in the *Open* state, and the original number of paths of the vpath is more than one. You can either use the Add Paths to Available Data Path Devices SMIT panel or run the **addpaths** command from the AIX command line. Go to “Dynamically adding paths to SDD vpath devices” on page 36 for more information about the **addpaths** command.

---

## Using LVM with SDD

This section describes how to use the Logical Volume Manager (LVM) with SDD to import and export volume groups, recover from mixed volume groups, extend an existing volume group, and back up and restore files belonging to a volume group.

### Importing volume groups with SDD

You can import a new volume group definition from a set of physical volumes with SDD vpath devices using the Volume Groups SMIT panel.

**Note:** To use this feature, you must either have root user authority or be a member of the system group.

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

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Import a Volume Group** and press **Enter**. The Import a Volume Group panel is displayed.
6. In the Import a Volume Group panel, perform the following tasks:
  - a. Enter the volume group that you want to import.
  - b. Enter the physical volume that you want to import.
  - c. Press **Enter** after making the changes.

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

## Exporting a volume group with SDD

You can export a volume group definition from the system with SDD vpath devices using the Volume Groups SMIT panel.

The **exportvg** command removes the definition of the volume group specified by the Volume Group parameter from the system. Because all system knowledge of the volume group and its contents are removed, an exported volume group is no longer accessible. The **exportvg** command does not modify any user data in the volume group.

A volume group is an unshared resource within the system; it should not be accessed by another system until it has been explicitly exported from its current system and imported on another. The primary use of the **exportvg** command, coupled with the **importvg** command, is to allow portable volumes to be exchanged between systems. Only a complete volume group can be exported, not individual physical volumes.

Using the **exportvg** command and the **importvg** command, you can also switch ownership of data on physical volumes shared between two systems.

**Note:** To use this feature, you must either have root user authority or be a member of the system group.

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

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Export a Volume Group** and press **Enter**. The Export a Volume Group panel is displayed.
6. Enter the volume group to export and press **Enter**.

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

## Recovering from mixed volume groups

When an SDD volume group is not active (that is, varied off), and certain AIX system administrative operations cause a device reconfiguration, a pvid attribute will be created for the supported storage device hdisks. This will cause the SDD volume group to become a mixed volume group. The following command is an example of a command that does this:

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

Run the **dpovgfix** shell script to recover a mixed volume group. The syntax is **dpovgfix vg-name**. The script searches for an SDD vpath device corresponding to each hdisk in the volume group and replaces the hdisk with the SDD vpath device. In order for the shell script to be run, all mounted file systems of this volume group have to be unmounted. After successful completion of the **dpovgfix** shell script, mount the file systems again.



## Extending an existing SDD volume group

You can extend a volume group with SDD vpath devices using the Logical Volume Groups SMIT panel. The SDD vpath devices to be added to the volume group should be chosen from those that can provide failover protection. It is possible to add an SDD vpath device to an SDD volume group that has only a single path (vpath0 on 63) and then add paths later by reconfiguring the supported storage device. With a single path, failover protection is not provided. (See “Dynamically adding paths to SDD vpath devices” on page 36 for information about adding paths to an SDD device.)

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

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Group** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Add a Data Path Volume to a Volume Group** and press **Enter**.
6. Enter the volume group name and physical volume name and press **Enter**. You can also use the F4 key to list all the available SDD devices, and you can select the devices that you want to add to the volume group.

If you use a script file to extend an existing SDD volume group, you must modify your script file and replace the **extendvg** command with the **extendvg4vp** command.

## Backing up all files belonging to an SDD volume group

You can back up all files belonging to a specified volume group with SDD vpath devices using the Volume Groups SMIT panel.

To back up a volume group with SDD devices, go to “Accessing the Backup a Volume Group with Data Path Devices SMIT panel” on page 75.

If you use a script file to back up all files belonging to a specified SDD volume group, you must modify your script file and replace the **savevg** command with the **savevg4vp** command.

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

## Restoring all files belonging to an SDD volume group

You can restore all files belonging to a specified volume group with SDD vpath devices using the Volume Groups SMIT panel.

To restore a volume group with SDD vpath devices, go to “Accessing the Remake a Volume Group with Data Path Devices SMIT panel” on page 75.

If you use a script file to restore all files belonging to a specified SDD volume group, you must modify your script file and replace the **restvg** command with the **restvg4vp** command.

## SDD-specific SMIT panels

SDD supports several special SMIT panels. Some SMIT panels provide SDD-specific functions, while other SMIT panels provide AIX functions (but require SDD-specific commands). For example, the Add a Volume Group with Data Path Devices function uses the SDD **mkvg4vp** command, instead of the AIX **mkvg** command. Table 11 lists the SDD-specific SMIT panels and how you can use them.

Table 11. SDD-specific SMIT panels and how to proceed

| SMIT panels                                   | How to proceed using SMITTY – Go to:                                               | Equivalent SDD command              |
|-----------------------------------------------|------------------------------------------------------------------------------------|-------------------------------------|
| Display Data Path Device Configuration        | “Accessing the Display Data Path Device Configuration SMIT panel” on page 72       | <b>lsvpcfg</b>                      |
| Display Data Path Device Status               | “Accessing the Display Data Path Device Status SMIT panel” on page 72              | <b>datapath query device</b>        |
| Display Data Path Device Adapter Status       | “Accessing the Display Data Path Device Adapter Status SMIT panel” on page 73      | <b>datapath query adapter</b>       |
| Define and Configure all Data Path Devices    | “Accessing the Define and Configure All Data Path Devices SMIT panel” on page 73   | <b>cfallvpath</b>                   |
| Add Paths to Available Data Path Devices      | “Accessing the Add Paths to Available Data Path Devices SMIT panel” on page 73     | <b>addpaths</b>                     |
| Configure a Defined Data Path Device          | “Accessing the Configure a Defined Data Path Device SMIT panel” on page 74         | <b>mkdev</b>                        |
| Remove a Data Path Device                     | “Accessing the Remove a Data Path Device SMIT panel” on page 74                    | <b>rmdev</b>                        |
| Add a Volume Group with Data Path Devices     | “Accessing the Add a Volume Group with Data Path Devices SMIT panel” on page 74    | <b>mkvg4vp</b>                      |
| Add a Data Path Volume to a Volume Group      | “Accessing the Add a Data Path Volume to a Volume Group SMIT panel” on page 74     | <b>extendvg4vp</b>                  |
| Remove a Physical Volume from a Volume Group  | “Accessing the Remove a Physical Volume from a Volume Group SMIT panel” on page 75 | <b>exportvg</b> <i>volume_group</i> |
| Back Up a Volume Group with Data Path Devices | “Accessing the Backup a Volume Group with Data Path Devices SMIT panel” on page 75 | <b>savevg4vp</b>                    |
| Remake a Volume Group with Data Path Devices  | “Accessing the Remake a Volume Group with Data Path Devices SMIT panel” on page 75 | <b>restvg</b>                       |

## Accessing the Display Data Path Device Configuration SMIT panel

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

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Configuration** and press **Enter**
5. The following example shows the Data Path Devices panel:

```
+-----+
| |
| Display Data Path Device Configuration |
| |
| Type or select values in entry fields. |
| Press Enter AFTER making all desired changes. |
| |
| Select Query Option [Entry Fields] |
| Device Name/ Device Model all devices +
| [] |
+-----+
```

The Select Query Option has three options:

### All devices

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

### Device name

This option runs **lsvpcfg <device name>** and only the specified device is displayed. Enter a device name in the Device Name/Device Model field.

### Device model

This option runs **lsvpcfg -d <device model>** and only devices with the specified device model are displayed. Enter a device model in the Device Name/Device Model field.

See “lsvpcfg” on page 77 for detailed information about the **lsvpcfg** command.

## Accessing the Display Data Path Device Status SMIT panel

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

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Status** and press **Enter**.
5. The following example shows the Data Path Devices Status panel:

```
+-----+
| |
| Display Data Path Device Status |
| |
| Type or select values in entry fields. |
| Press Enter AFTER making all desired changes. |
| |
| [Entry Fields] |
+-----+
```

|                             |             |   |
|-----------------------------|-------------|---|
| Select Query Option         | all devices | + |
| Device Number/ Device Model | [ ]         |   |

---

The Select Query Option has 3 options:

**All devices**

This option runs **datapath query device** and all the data path devices are displayed. No entry is required in the Device Name/Device Model field.

**Device number**

This option runs **datapath query device <device number>** and only the specified device is displayed. Enter a device number in the Device Name/Device Model field.

**Device model**

This option runs **datapath query device -d <device model>** and only devices with the specified device model are displayed. Enter a device model in the Device Name/Device Model field.

See “datapath query device” in Chapter 13 of the *IBM System Storage Multipath Subsystem Device Driver User’s Guide* for detailed information about the **datapath query device** command.

## Accessing the Display Data Path Device Adapter Status SMIT panel

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

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

Perform the following steps to access the Define and Configure All Data Path Devices panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Define and Configure All Data Path Devices** and press **Enter**.

## Accessing the Add Paths to Available Data Path Devices SMIT panel

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

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.

4. Select **Add Paths to Available Data Path Devices** and press **Enter**.

## Accessing the Configure a Defined Data Path Device SMIT panel

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

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Configure a Defined Data Path Device** and press **Enter**.

## Accessing the Remove a Data Path Device SMIT panel

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

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Remove a Data Path Device** and press **Enter**.

## Accessing the Add a Volume Group with Data Path Devices SMIT panel

Perform the following steps to access the Add a volume group with data path devices panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Add Volume Group with Data Path Devices** and press **Enter**.

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

## Accessing the Add a Data Path Volume to a Volume Group SMIT panel

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

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical)** and press **Enter**. The System Storage Management (Physical & Logical) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Group** and press **Enter**. The Volume Group panel is displayed.
5. Select **Add a Data Path Volume to a Volume Group** and press **Enter**.

6. Enter the volume group name and physical volume name and press **Enter**. Alternately, you can use the F4 key to list all the available SDD vpath devices and use the F7 key to select the physical volumes that you want to add.

## Accessing the Remove a Physical Volume from a Volume Group SMIT panel

Perform the following steps to access the Remove a Physical Volume from a Volume Group panel:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
3. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
4. Select **Set Characteristics of a Volume Group** and press **Enter**. The Set Characteristics of a Volume Group panel is displayed.
5. Select **Remove a Physical Volume from a Volume Group** and press **Enter**. The Remove a Physical Volume from a Volume Group panel is displayed.

## Accessing the Backup a Volume Group with Data Path Devices SMIT panel

Perform the following steps to access the Back Up a Volume Group with Data Path Devices panel and to backup a volume group with SDD devices:

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Back Up a Volume Group with Data Path Devices** and press **Enter**. The Back Up a Volume Group with Data Path Devices panel is displayed.
6. In the Back Up a Volume Group with Data Path Devices panel, perform the following steps:
  - a. Enter the Backup DEVICE or FILE name.
  - b. Enter the Volume Group to backup.
  - c. Press **Enter** after making all required changes.

**Tip:** You can also use the F4 key to list all the available SDD devices, and you can select the devices or files that you want to backup.

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

## Accessing the Remake a Volume Group with Data Path Devices SMIT panel

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

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

2. Select **System Storage Management (Physical & Logical Storage)** and press **Enter**. The System Storage Management (Physical & Logical Storage) panel is displayed.
3. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
4. Select **Volume Groups** and press **Enter**. The Volume Groups panel is displayed.
5. Select **Remake a Volume Group with Data Path Devices** and press **Enter**. The Remake a Volume Group with Data Path Devices panel is displayed.
6. Enter the Restore DEVICE or FILE name that you want to restore, and press **Enter**. You can also press **F4** to list all the available SDD devices, and you can select the devices or files that you want to restore.

---

## SDD utility programs

The following SDD utility programs are available:

### addpaths

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

This command will open a new path (or multiple paths) automatically if the SDD vpath device is in *Open* state. You can either use the Add Paths to Available Data Path Devices SMIT panel or run the **addpaths** command from the AIX command line.

The syntax for this command is:

```
▶▶—addpaths—◀◀
```

For more information about this command, go to “Dynamically adding paths to SDD vpath devices” on page 36.

### hd2vp and vp2hd

The SDD provides two conversion scripts, **hd2vp** and **vp2hd**. The **hd2vp** script converts a volume group from supported storage device hdisks to the SDD vpath devices, and the **vp2hd** script converts a volume group from the SDD vpath devices to supported storage device hdisks. Use the **vp2hd** program when you want to configure your applications back to original supported storage device hdisks, or when you want to remove the SDD from your AIX host system.

The syntax for these conversion scripts is as follows:

```
▶▶—hd2vpvgname—◀◀
```

```
▶▶—vp2hdvgname—◀◀
```



*vgroupname*

Specifies the volume group name to be converted.

## dpovgfix

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

Performing AIX system management operations on adapters and `hdisk` devices can cause original supported storage device `hdisks` to be contained within an SDD volume group. This is known as a mixed volume group. Mixed volume groups happen when an SDD volume group is not active (varied off), and certain AIX commands to the `hdisk` put the `pvid` attribute of `hdisk` back into the ODM database. The following is an example of a command that does this:

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

If this disk is an active `hdisk` of an SDD `vpath` device that belongs to an SDD volume group, and you run the `varyonvg` command to activate this SDD volume group, LVM might pick up the `hdisk` device instead of the SDD `vpath` device. The result is that an SDD volume group partially uses the SDD `vpath` devices, and partially uses supported storage device `hdisk` devices. This causes the volume group to lose path-failover capability for that physical volume. The `dpovgfix` script tool fixes this problem.

The syntax for this command is:

```
▶▶—dpovgfix vgroupname—————▶▶
```

*vgroupname*

Specifies the volume group name of the mixed volume group to be recovered.

## lsvpcfg

You can use the `lsvpcfg` script tool to display the configuration state of SDD devices. This displays the configuration state for all SDD devices. The `lsvpcfg` command can be issued in three ways.

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

```
lsvpcfg
```

See “Verifying the SDD configuration” on page 35 for an example of the output and what it means.

2. The command can also be issued using the SDD `vpath` device name as a parameter. The syntax for this command is:

```
lsvpcfg 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 35 for an explanation of the output.

- The command can also be issued using the device model as a parameter. The option to specify a device model cannot be used when you specify an SDD vpath device. The syntax for this command is:

```
lsvpcfg device model
```

The following are examples of valid device models:

**2105** Display all 2105 models (ESS).

**2105F** Display all 2105 F models (ESS).

**2105800**

Display all 2105 800 models (ESS).

**2145** Display all 2145 models (SAN Volume Controller).

**2107** Display all DS8000 models.

**1750** Display all DS6000 models.

## mkvg4vp

You can use the **mkvg4vp** command to create an SDD volume group. For more information about this command, go to “Configuring volume groups for failover protection” on page 63. For information about the flags and parameters for this command, go to:

[http://publib16.boulder.ibm.com/doc\\_link/en\\_US/a\\_doc\\_lib/cmds/aixcmds3/mkvg.htm](http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/cmds/aixcmds3/mkvg.htm)

The syntax for this command is:

```

▶▶ mkvg4vp [-S*] [-d MaxPVs] [-B] [-G] [-f] [-q **]
▶ [-C | -c [-x]] [-i] [-s PPSize] [-n]
▶ [-m MaxPVsize | -t factor] [-V MajorNumber] [-L LTGsize***]
▶ [-y VGname] PVname

```

\* for AIX 5.3 and later only

\*\* for AIX 5.2 and later only

\*\*\* for AIX 5.1 and later only

## extendvg4vp

You can use the **extendvg4vp** command to extend an existing SDD volume group. For more information about this command, go to “Extending an existing SDD volume group” on page 70. For information about the flag and parameters for this command, go to:

[http://publib16.boulder.ibm.com/doc\\_link/en\\_US/a\\_doc\\_lib/cmds/aixcmds2/extendvg.htm](http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/cmds/aixcmds2/extendvg.htm)

The syntax for this command is:

```
▶▶ extendvg4vp -f VGname—PVname ▶▶▶
```

## querysn

You can use the **querysn** command to exclude supported storage device (hdisk) from the SDD vpath configuration. You must run this command before the SDD vpath devices are configured. The **querysn** command saves the serial number for the logical device (hdisk) to an exclude file (/etc/vpexclude). During the SDD configuration, the SDD configure methods read all serial numbers listed in this file and exclude these devices from the SDD configuration.

The syntax for this command is:

```
▶▶ querysn -d -l device-name ▶▶▶
```

- l Specifies the logical number of the supported storage device (hdiskN). This is not the SDD device name.
- d When this optional flag is set, the **querysn** command deletes all existing contents from this exclude file and then writes the new serial number into the file; otherwise, it appends the new serial number at the end of the file.

### device name

Specifies the supported storage device (hdiskN).

### Example:

```
querysn -l hdisk10
```

### Notes:

1. Do not use the **querysn** command to exclude a device if you want the device to be configured by SDD.
2. If the supported storage device LUN has multiple configurations on a server, use the **querysn** command on only one of the logical names of that LUN.
3. You should not use the **querysn** command on the same logical device multiple times. Using the **querysn** command on the same logical device multiple times results in duplicate entries in the /etc/vpexclude file, and the system administrator will have to administer the file and its content.
4. Issue the **querysn** command with the -d flag to delete all existing contents from the exclude file and then write the new serial number into the file. If you want to remove only one device from the /etc/vpexclude file, edit the /etc/vpexclude with the vi editor and delete the line containing the device name. To replace a manually excluded device in the SDD configuration, open the /etc/vpexclude file with a text editor (for example, vi) and delete the line containing the device name. For detailed instructions on the proper procedure, see “Replacing manually excluded devices in the SDD configuration” on page 41.

## lquerypr

See “Persistent reserve command tool” on page 80.

## sddgetdata

See Appendix A of the *IBM System Storage Multipath Subsystem Device Driver User's Guide*, which describes the use of sddgetdata to collect information for problem determination.

---

## Persistent reserve command tool

SDD offers the **lquerypr** tool to provide a set of persistent reserve functions. This tool supports the following persistent reserve service actions:

- Read persistent reservation key
- Release persistent reserve
- Preempt-abort persistent reserve
- Clear persistent reserve and registrations

**Note:** Caution must be taken with the command, especially when implementing preempt-abort or clear persistent reserve service action. With preempt-abort service action not only the current persistent reserve key is preempted; it also aborts tasks on the LUN that originated from the initiators that are registered with the preempted key. With clear service action, both persistent reservation and reservation key registrations are cleared from the device or LUN.

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

### lquerypr command

#### Purpose

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

#### Syntax

```
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, such as read reservation key, release persistent reservation, preempt-abort persistent reservation, and clear persistent reservation.

**Note:** This command can only be used when the device is not already opened.

#### Flags

- p** If the persistent reservation key on the device is different from the current host reservation key, it preempts the persistent reservation key on the device.

- c If there is a persistent reservation key on the device, it removes any persistent reservation and clears all reservation key registration on the device.
- r Removes the persistent reservation key on the device made by this host.
- v Displays the persistent reservation key if it exists on the device.
- V Verbose mode. Prints detailed message.

#### Return code

If the command is issued without options of -p, -r, or -c, the command will return 0 under two circumstances.

1. There is no persistent reservation key on the device.
2. The device is reserved by the current host.

If the persistent reservation key is different from the host reservation key, the command will return 1. If the command fails, it returns 2. If the device is already opened on a current host, the command returns 3.

#### Example

1. To query the persistent reservation on a device, enter **lquerypr -h/dev/vpath30**.  
This command queries the persistent reservation on the device without displaying. If there is a persistent reserve on a disk, it returns 0 if the device is reserved by the current host. It returns 1 if the device is reserved by another host.
2. To query and display the persistent reservation on a device, enter **lquerypr -vh/dev/vpath30**.  
Same as Example 1. In addition, it displays the persistent reservation key.
3. To release the persistent reservation if the device is reserved by the current host, enter **lquerypr -rh/dev/vpath30**.  
This command releases the persistent reserve if the device is reserved by the current host. It returns 0 if the command succeeds or the device is not reserved. It returns 2 if the command fails.
4. To reset any persistent reserve and clear all reservation key registrations, enter **lquerypr -ch/dev/vpath30**.  
This command resets any persistent reserve and clears all reservation key registrations on a device. It returns 0 if the command succeeds, or 2 if the command fails.
5. To remove the persistent reservation if the device is reserved by another host, enter **lquerypr -ph/dev/vpath30**.  
This command removes an existing registration and persistent reserve from another host. It returns 0 if the command succeeds or if the device is not persistent reserved. It returns 2 if the command fails.

---

## Using supported storage devices directly

When you configure your host adapters, the AIX disk driver creates hdisks for the supported storage device LUNs. When you configure SDD, it creates SDD vpath devices for supported storage device LUNs. Therefore, with SDD installed on a system, an application now has two ways to access the supported storage devices.

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

Applications can access SDD vpath devices either by using raw devices or logical devices. If the application is accessing SDD vpath devices using logical devices, a volume group must be created with SDD vpath devices.

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

1. Enter **smitty** from your desktop window. SMIT is displayed.
2. Select **Devices** and press **Enter**. The Devices panel is displayed.
3. Select **Data Path Devices** and press **Enter**. The Data Path Devices panel is displayed.
4. Select **Display Data Path Device Configuration** and press **Enter**.
5. To display all SDD vpath devices with their attached multiple paths (hdisks), select all SDD vpath devices for **Select Query Option**, leave Device Name/Device Model blank, and press **Enter**.
6. Search the list of hdisks to locate the hdisks that your application is using.
7. Replace each hdisk with its corresponding SDD vpath device.

**Note:** Depending upon your application, the manner in which you replace these files is different. If this is a new application, use the SDD vpath device rather than hdisk to use the SDD load-balancing and failover features.

**Note:** Alternately, you can enter **lsvpcfg** from the command-line interface rather than using SMIT. This displays all configured SDD vpath devices and their underlying paths (hdisks).

---

## Using supported storage devices through AIX LVM

If your application accesses supported storage devices through LVM, determine that the physical volumes of the volume group that the application is accessing are SDD-supported storage devices. Then perform the following steps to convert the volume group from the original supported storage device hdisks to the SDD vpath devices:

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

- b. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
  - c. Select **File Systems** and press **Enter**. The File Systems panel is displayed.
  - d. Select **List All File Systems** to locate all file systems that have the same mount points as the logical volumes and press **Enter**. The file systems are listed.
  - e. Note the file system name of that volume group and the file system mount point, if it is mounted.
  - f. Unmount these file systems.
8. Enter the following command to convert the volume group from the supported storage device hdisks to SDD multipath vpath devices:
 

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

When the conversion is complete, your application now accesses supported storage device physical LUNs through SDD vpath devices. This provides load-balancing and failover protection for your application.

---

## Migrating a non-SDD volume group to a supported storage device SDD multipath volume group in concurrent mode

Before you migrate your non-SDD volume group to an SDD volume group, make sure that you have completed the following tasks:

- The SDD for the AIX host system is installed and configured. See “Verifying the currently installed version of SDD for SDD 1.3.3.11 (or earlier)” on page 25 or “Verifying the currently installed version of SDD for SDD 1.4.0.0 (or later)” on page 26.
- The supported storage devices to which you want to migrate have multiple paths configured per LUN. To check the state of your SDD configuration, use the System Management Interface Tool (SMIT) or issue the **lsvpcfg** command from the command line. To use SMIT:
  - Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
  - Select **Devices** and press **Enter**. The Devices panel is displayed.
  - Select **Data Path Device** and press **Enter**. The Data Path Device panel is displayed.
  - Select **Display Data Path Device Configuration** and press **Enter**. A list of the SDD vpath devices and whether there are multiple paths configured for the devices is displayed.
- Ensure that the SDD vpath devices that you are going to migrate to do not belong to any other volume group, and that the corresponding physical device (supported storage device LUN) does not have a pvid written on it. Enter the **lsvpcfg** command output to check the SDD vpath devices that you are going to use for migration. Make sure that there is no pv displayed for this SDD vpath device and its paths (hdisks). If a LUN has never belonged to any volume group, there is no pvid written on it. In case there is a pvid written on the LUN and the LUN does not belong to any volume group, you need to clear the pvid from the LUN before using it to migrate a volume group. The commands to clear the pvid are:

```
chdev -l hdiskW -a pv=clear
chdev -l vpathW -a pv=clear
```



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

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

1. Add new SDD vpath devices to an existing non-SDD volume group:
  - a. Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
  - b. Select **System Storage Management (Physical & Logical)** and press **Enter**. The System Storage Management (Physical & Logical) panel is displayed.
  - c. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
  - d. Select **Volume Group** and press **Enter**. The Volume Group panel is displayed.
  - e. Select **Add a Data Path Volume to a Volume Group** and press **Enter**.
  - f. Enter the volume group name and physical volume name and press **Enter**. Alternately, you can use the F4 key to list all the available SDD vpath devices and use the F7 key to select the physical volumes that you want to add.
2. Enter the **smitty mklvcopy** command to mirror logical volumes from the original volume to an SDD supported storage device volume. Use the new SDD vpath devices for copying all logical volumes. Do not forget to include JFS log volumes.

**Note:** The command **smitty mklvcopy** copies one logical volume at a time. A fast-path command to mirror *all* the logical volumes on a volume group is **mirrorvg**.

3. Synchronize logical volumes (LVs) or force synchronization. Enter the **smitty syncvg** command to synchronize all the volumes:

There are two options on the smitty panel:

- Synchronize by Logical Volume
- Synchronize by Physical Volume

The fast way to synchronize logical volumes is to select the **Synchronize by Physical Volume** option.

4. Remove the mirror and delete the original LVs. Enter the **smitty rmlvcopy** command to remove the original copy of the logical volumes from all original non-SDD physical volumes.
5. Enter the **smitty reducevg** command to remove the original non-SDD vpath devices from the volume group. The Remove a Physical Volume panel is displayed. Remove all non-SDD devices.

**Note:** A non-SDD volume group refers to a volume group that consists of non-supported storage devices or supported storage hdisk devices.

## Detailed instructions for migrating a non-SDD volume group to a supported storage device SDD multipath volume group in concurrent mode

This procedure shows how to migrate an existing AIX volume group to use SDD vpath devices that have multipath capability. You do not take the volume group out of service. The example shown starts with a volume group, `vg1`, made up of one supported storage device, `hdisk13`.

To perform the migration, you must have SDD vpath devices available that are greater than or equal to the size of each of the `hdisk`s making up the volume group. In this example, there is an SDD device, `vpath12`, with two paths, `hdisk14` and `hdisk30`, to which we will migrate the volume group.

1. Add the SDD vpath device to the volume group as an Available volume:
  - a. Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
  - b. Select **System Storage Management (Physical & Logical)** and press **Enter**. The System Storage Management (Physical & Logical) panel is displayed.
  - c. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
  - d. Select **Volume Group** and press **Enter**. The Volume Group panel is displayed.
  - e. Select **Add a Data Path Volume to a Volume Group** and press **Enter**.
  - f. Enter **vg1** in the **Volume Group Name** field and enter **vpath12** in the **Physical Volume Name** field. Press **Enter**.  
You can also use the **extendvg4vp -f vg1 vpath12** command.
2. Mirror logical volumes from the original volume to the new SDD vpath device volume:
  - a. Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
  - b. Select **System Storage Management (Physical & Logical)** and press **Enter**. The System Storage Management (Physical & Logical) panel is displayed.
  - c. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
  - d. Select **Volume Group** and press **Enter**. The Volume Group panel is displayed.
  - e. Select **Mirror a Volume Group** and press **Enter**. The Mirror a Volume Group panel is displayed.
  - f. Enter a volume group name and a physical volume name. Press **Enter**.  
You can also enter the **mirrorvg vg1 vpath12** command.
3. Synchronize the logical volumes in the volume group:
  - a. Enter **smitty** and press **Enter** from your desktop window. The System Management Interface Tool panel is displayed.
  - b. Select **System Storage Management (Physical & Logical)** and press **Enter**. The System Storage Management (Physical & Logical) panel is displayed.
  - c. Select **Logical Volume Manager** and press **Enter**. The Logical Volume Manager panel is displayed.
  - d. Select **Volume Group** and press **Enter**. The Volume Group panel is displayed.

- e. Select **Synchronize LVM Mirrors** and press **Enter**. The Synchronize LVM Mirrors panel is displayed.
- f. Select **Synchronize by Physical Volume**.

You can also enter the `syncvg -p hdisk13 vpath12` command.

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

You can also enter the command:

```
rm1vcopy loglv01 1 hdisk13
rm1vcopy lv01 1 hdisk13
```

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

You can also enter the `reducevg vg1 hdisk13` command.

---

## Using the trace function

SDD supports AIX trace functions. The SDD trace ID is 2F8. Trace ID 2F8 traces routine entry, exit, and error paths of the algorithm. To use it, manually turn on the trace function before the program starts to run, then turn off the trace function either after the program stops, or any time you need to read the trace report.

By default, SDD disables the AIX trace functions. To enable the AIX trace functions for SDD, follow these steps:

1. Enter `pathctest -d` . (for example, `pathctest -d 0`)
2. Enter `777`.
3. Enter `20` to open the device.
4. Enter `3` (as option `NO_DELAY`).
5. Enter `90` (enable or disable the AIX trace). Follow the prompt: enter `1` to enable.

Then you can start the trace function.

To start the trace function, enter:

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

To stop the trace function, enter:

**trcstop**

To read the report, enter:

**trcrpt | pg**

To save the trace data to a file, enter:

**trcrpt > filename**

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



---

## Chapter 2. Using SDDPCM on an AIX host system

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

This chapter provides a general view of the SDDPCM path control module, including where it resides on the I/O stack in operating system and the features and functions that it supports. This chapter also provides procedures for the following management tasks:

- Install SDDPCM on an AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1 host system
- Configure SDDPCM MPIO-capable devices
- Uninstall the SDDPCM module on an AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1 host system
- Migrate disk storage MPIO-capable devices from the AIX default PCM to SDDPCM
- Migrate disk storage MPIO-capable devices from SDDPCM to the AIX default PCM or SDD
- Installing SDDPCM from AIX NIM SPOT to client SAN boot disk or internal boot disk

Figure 1 on page 90 shows the position of SDDPCM in the protocol stack. I/O operations are sent to the AIX disk driver. The SDDPCM path selection routine is invoked to select an appropriate path for each I/O operation.

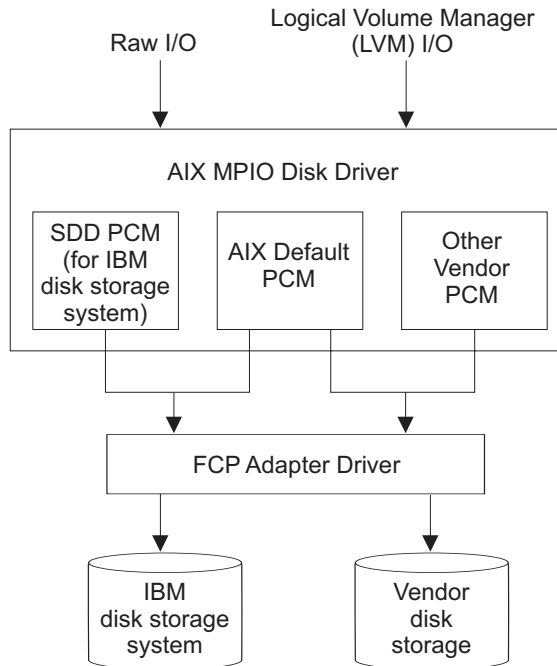


Figure 1. SDDPCM in the protocol stack

For detailed information about MPIO support on AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1, visit the following Web site:

[http://publib16.boulder.ibm.com/pseries/en\\_US/aixbman/baseadm/managed\\_mpio.htm](http://publib16.boulder.ibm.com/pseries/en_US/aixbman/baseadm/managed_mpio.htm)

AIX MPIO-capable device drivers will automatically discover, configure and make available every storage device path. SDDPCM manages the paths to provide:

- High availability and load balancing of storage I/O
- Automatic path-failover protection
- Concurrent download of supported storage devices licensed machine code
- Prevention of a single-point-failure

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

[www.ibm.com/servers/storage/support/software/sdd](http://www.ibm.com/servers/storage/support/software/sdd)

SDD and SDDPCM are exclusive software packages on a server. You cannot install both software packages on a server for supported storage devices. When supported storage devices are configured as non-MPIO-capable devices (that is, multiple logical device instances are created for a physical LUN), you should install SDD to get multipath support.

You must install SDDPCM in order to configure supported storage devices into MPIO-capable-devices (where only one logical device instance is created for a physical LUN). In order to run SDDPCM on AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1, you must install all of the latest PTFs for that OS level.

Configuring supported storage devices into MPIO-capable devices depends on the version of supported storage devices host attachment installed on the system.



To configure supported storage devices as non-MPIO-capable devices, install `ibm2105.rte` (version 32.6.100.x) or `devices.fcp.disk.ibm.rte` (version 1.0.0.x), or both. To configure supported storage devices as MPIO-capable devices, install `devices.fcp.disk.ibm.mpio.rte` (version of 1.0.0.11).

For the latest version of the supported storage devices host attachment package, refer to the Readme file on the SDD download Web site:

[www.ibm.com/servers/storage/support/software/sdd](http://www.ibm.com/servers/storage/support/software/sdd)

**Note:** SDDPCM does not support SCSI storage devices.

With SDD 1.6.0.0 (or later), SDDPCM and SDD cannot coexist on a AIX server. If a server connects to any supported storage devices, all devices must be configured either as non-MPIO-capable devices or as MPIO-capable devices.

---

## Supported SDDPCM features

The following SDDPCM features are supported in this release:

- 32- and 64-bit kernels
- Four types of reserve policies:
  - No\_reserve policy
  - Exclusive host access single path policy
  - Persistent reserve exclusive host policy
  - Persistent reserve shared host access policy
- Three path-selection algorithms:
  - Failover
  - Round robin
  - Load balancing
- Automatic failed paths reclamation by healthchecker
- Failback error-recovery algorithm
- Fibre-channel dynamic device tracking
- Support for all ESS FCP, DS8000, DS6000, and SAN Volume Controller devices
- Support SAN boot device on MPIO supported storage devices

**Note:** SDDPCM supports ESS devices as SAN boot devices, starting from AIX52 TL06 and AIX53 TL02. SDDPCM supports DS8000, DS6000, and SAN Volume Controller devices as SAN boot devices, starting from AIX52 TL07 and AIX53 TL03.

- Support for external supported MPIO storage devices as the primary or secondary dump device
- Support storage devices multipath devices as system paging space
- Support SDDPCM server daemon enhanced path health check function
- Support a maximum of 1200 LUNs
- Dynamically adding paths or adapters
- Dynamically removing paths or adapters
- Dynamically changing the device path selection algorithm
- Dynamically changing the device `hc_interval`
- Dynamically changing the device `hc_mode`

- Web-based System Manager (WebSM) for MPIO supported storage devices (Refer to [www-1.ibm.com/servers/aix/wsm/](http://www-1.ibm.com/servers/aix/wsm/) for more information about WebSM.)
- Reserve last path of a device in OPEN mode
- Support the essutil Product Engineering tool in SDDPCM's **pcmpath** command line program
- Support HACMP with Enhanced Concurrent Mode volume group in concurrent resource groups and nonconcurrent resource groups
- Support GPFS in AIX 5.2 TL06 (or later), 5.3 TL02 (or later), and AIX 6.1.
- Support Virtual I/O server with AIX 5.3 or later and AIX 6.1.

---

## Unsupported SDDPCM features

The following SDDPCM features are not currently supported.

- HACMP with persistent reservation policies
- Multiple hosts sharing supported storage MPIO devices without supported clustering software installed
- HACMP with non-Enhanced Concurrent Mode volume group
- Mirror Write Consistency with volume group in a concurrent resource group

---

## Verifying the hardware and software requirements

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

### Hardware

The following hardware components are needed:

- Supported storage devices (FCP devices only)
- One or more switches, if the supported storage devices is not direct-attached
- Host system
- Fibre-channel adapters and cables

### Software

The following software components are needed:

- AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1 operating system.
- Fibre-channel device drivers
- One of the following installation packages:
  - `devices.sddpcm.52.rte` (version 2.2.0.0 or later)
  - `devices.sddpcm.53.rte` (version 2.2.0.0 or later)
  - `devices.sddpcm.61.rte` (version 2.2.0.0 or later)
- Supported storage devices: `devices.fcp.disk.ibm.mpio.rte` (version of 1.0.0.11 or later) host attachment package for SDDPCM

## Unsupported environments

SDDPCM does not support:

- ESS SCSI devices
- A host system with both a SCSI and fibre-channel connection to a shared ESS logical unit number (LUN)

- Single-path mode during code distribution and activation of LMC nor during any supported storage devices concurrent maintenance that impacts the path attachment, such as a supported storage device host-bay-adapter replacement.

## Host system requirements

To successfully install SDDPCM for supported storage devices, you must have AIX 5.2 TL08 (or later), AIX 5.3 TL04 (or later), or AIX 6.1 installed on your host system along with the AIX required fixes, APARs, and microcode updates that are identified on the following Web site:

[www-1.ibm.com/servers/storage/support/](http://www-1.ibm.com/servers/storage/support/)

## Supported storage device requirements

To successfully install SDDPCM, ensure that the `devices.fcp.disk.ibm.mpio.rte` (version 1.0.0.11) package for supported storage device attachment is installed on the server.

## Fibre requirements

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

[www-1.ibm.com/servers/eserver/support/](http://www-1.ibm.com/servers/eserver/support/)

If your host has only one fibre-channel adapter, it requires you to connect through a switch to multiple supported storage device ports. You should have at least two fibre-channel adapters to prevent data loss due to adapter hardware failure or software failure.

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

[www.ibm.com/servers/storage/support](http://www.ibm.com/servers/storage/support)

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

- The AIX host system is an IBM RS/6000 or IBM System p with AIX 5.2 TL07 (or later), AIX 5.3 TL04 (or later), or AIX 6.1.
- The AIX host system has the fibre-channel device drivers installed along with all latest APARs.
- The host system can be a single processor or a multiprocessor system, such as SMP.
- A fiber-optic cable connects each fibre-channel adapter to a supported storage system port.
- If you need the SDDPCM I/O load-balancing and failover features, ensure that a minimum of two paths to a device are attached.

---

## Preparing for SDDPCM installation

The SDDPCM installation package installs a number of major files on your AIX system. The major files that are part of the SDDPCM installation package are:

| File name              | Description                                                                                   |
|------------------------|-----------------------------------------------------------------------------------------------|
| <code>sddpcmrtl</code> | A dynamically loaded module added to the device configuration methods to extend the supported |

|                           |                                                                                                                                                   |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
|                           | storage devices device configuration methods to facilitate the configuration operations of the PCM KE                                             |
| <b>sddpcmke</b>           | A dynamically-loaded module added to the AIX 5L kernel that provides path management functions for supported storage devices                      |
| <b>sdduserke</b>          | A dynamically-loaded module added to the AIX 5L kernel that provides the API to sddpcmke                                                          |
| <b>pcmpath</b>            | SDDPCM command line tool                                                                                                                          |
| <b>pcmsrv</b>             | Daemon for enhanced path healthcheck, and First Time Data Capture                                                                                 |
| <b>sample_pcmsrv.conf</b> | The sample SDDPCM server daemon configuration file                                                                                                |
| <b>fcppcmmap</b>          | Collects supported storage devices fibre-channel device information through SCSI commands                                                         |
| <b>pcmquerypr</b>         | SDDPCM persistent reserve command tool                                                                                                            |
| <b>pcmggenprkey</b>       | SDDPCM persistent reserve command tool to generate persistent reserve key                                                                         |
| <b>relbootrsv</b>         | Release SCSI-2 reserve on boot devices or on active nonboot volume groups                                                                         |
| <b>sddpcmgetdata</b>      | Script to collect SDDPCM information, trace log files, and system error logs into an sddpcmdata_host_date_time.tar file for problem determination |

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

## Preparing for SDDPCM installation for supported storage devices

Before you install SDDPCM, you must complete the following tasks:

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

For more information about how to configure your supported storage device, refer to the Planning Guide for your product.

Before you install SDDPCM, you must:

- Determine that you have the correct installation package
- Remove the SDD package, if it is installed.
- Remove the ibm2105.rte (version 32.6.100.x) and/or devices. fcp.disk.ibm.rte (version 1.0.0.x), if they are installed.
- Install the AIX fibre-channel device drivers, if necessary.
- Verify and upgrade the fibre channel adapter firmware level

- Install the SDDPCM Host Attachment: `devices.fcp.disk.ibm.mpio.rte` ( version 1.0.0.11)

### **Determining the correct installation package**

SDDPCM can be installed only on an AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1 operating system. The package name of SDDPCM is the following:

- `devices.sddpcm.52.rte` for AIX 5.2
- `devices.sddpcm.53.rte` for AIX 5.3
- `devices.sddpcm.61.rte` for AIX 6.1

### **Determining if the SDD package is installed**

To determine if the SDD is installed:

1. Use the `lspp -l *ibmSdd*` and `lspp -l devices.sdd*` commands to determine if any SDD package is installed on the system.
2. If SDD is installed for supported storage device configuration, you must unconfigure and remove all SDD vpath devices, and then uninstall the SDD package. See “Removing SDD from an AIX host system” on page 38.

### **Determining if the `ibm2105.rte` package is installed**

To determine if the `ibm2105.rte` package is installed:

1. Use the `lspp -l *ibm2105*` command to determine if any `ibm2105.rte` with VRMF 32.6.100.XX is installed.
2. If `ibm2105.rte` is installed, you must:
  - a. Unconfigure and remove all supported storage devices.
  - b. Use `smitty` to uninstall the `ibm2105.rte` package. If `ibm2105.rte` is installed, you must remove all supported storage devices and uninstall the `ibm2105.rte` package. See “Removing SDD from an AIX host system” on page 38.

### **Determining if the `devices.fcp.disk.ibm.rte` package is installed**

To determine if the `devices.fcp.disk.ibm.rte` package is installed:

1. Use the `lspp -l devices.fcp.disk.ibm*` command to determine if any `devices.fcp.disk.ibm.rte` with VRMF 1.0.0.X is installed.
2. If `devices.fcp.disk.ibm.rte` is installed, you must:
  - a. Unconfigure and remove all supported storage devices.
  - b. Use `smitty` to uninstall the `devices.fcp.disk.ibm.rte` package. If `devices.fcp.disk.ibm.rte` is installed, you must remove all supported storage devices and uninstall the `devices.fcp.disk.ibm.rte` package. See “Removing SDD from an AIX host system” on page 38.

### **Installing the AIX fibre-channel device drivers**

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

[www-1.ibm.com/servers/storage/support/](http://www-1.ibm.com/servers/storage/support/)

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

1. Log in as the root user.
2. Load the compact disc into the CD-ROM drive.

3. From your desktop window, enter **smitty install\_update** and press **Enter** to go directly to the installation panels. The **Install and Update Software** menu is displayed.
4. Highlight **Install Software** and press **Enter**.
5. Press **F4** to display the **INPUT Device/Directory for Software** panel.
6. Select the compact disc drive that you are using for the installation; for example, **/dev/cd0**, and press **Enter**.
7. Press **Enter** again. The **Install Software** panel is displayed.
8. Highlight **Software to Install** and press **F4**. The **Software to Install** panel is displayed.
9. The fibre-channel device drivers include the following installation packages:

**devices.pci.df1080f9**

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

**devices.pci.df1000f9**

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

**devices.pci.df1000f7**

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

**devices.common.IBM.fc**

The FCP protocol driver.

**devices.fcp.disk**

The FCP disk driver.

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

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

```

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

```

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

```
instfix -iv | grep IYnnnnn
```

where *nnnnn* represents the APAR numbers.

If the APARs are listed, that means that they are installed. If they are installed, go to “Configuring supported storage MPIO-capable devices” on page 106. Otherwise, go to step 3.

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

## Verifying and upgrading the fibre channel adapter firmware level

Use the following procedures to verify and upgrade your current fibre channel adapter firmware level.

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

<http://www14.software.ibm.com/webapp/set2/firmware/gjsn>

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

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

(ZB).....S2F3.30X1

To verify the firmware level, ignore the second character in the **ZB** field. In the example, the firmware level is sf330X1.

3. If the adapter firmware level is at the latest level, there is no need to upgrade; otherwise, the firmware level must be upgraded. To upgrade the firmware level, go to “Upgrading the adapter firmware level.”

**Upgrading the adapter firmware level:** Upgrading the firmware level consists of downloading the firmware (microcode) from your AIX host system to the adapter. Before you upgrade the firmware, ensure that you have configured any fibre-channel-attached devices (see “Configuring fibre-channel-attached devices” on page 8). After the devices are configured, download the firmware from the AIX host system to the FCP adapter by performing the following steps:

1. Verify that the correct level of firmware is installed on your AIX host system. Go to the `/etc/microcode` directory and locate the file called `df1000f7.XXXXXXX` for feature code 6227 and `df1000f9.XXXXXXX` for feature code 6228, where `XXXXXX` is the level of the microcode. This file was copied into the `/etc/microcode` directory during the installation of the fibre-channel device drivers.
2. From the AIX command prompt, enter **diag** and press **Enter**.
3. Highlight the **Task Selection** option.
4. Highlight the **Download Microcode** option.
5. Press **Enter** to select all the fibre-channel adapters to which you want to download firmware. Press **F7**. The Download panel is displayed with one of the selected adapters highlighted. Press **Enter** to continue.
6. Highlight `/etc/microcode` and press **Enter**.
7. Follow the instructions that are displayed to download the firmware, one adapter at a time.

---

## Installing the AIX SDDPCM host attachment

The host attachment for SDDPCM adds 2105, 2145, 1750, or 2107 device information to allow AIX to properly configure 2105, 2145, 1750, or 2107 as MPIO-capable hdisks. The 2105, 2145, 1750, or 2107 device information allows AIX to perform the following actions:



- Identify the hdisks as 2105, 2145, 1750, or 2107 hdisks.
- Set default hdisk attributes, such as `queue_depth` and `timeout` values.
- Indicate to the AIX device driver configure method to set 2105, 2145, 1750, or a 2107 hdisk as MPIO devices.

The AIX SDDPCM Host Attachment package name is `devices.fcp.disk.ibm.mpio.rte`. See the SDDPCM Readme file for the latest Host Attachment level for the SDDPCM version that you are installing.

See the *Host Systems Attachment Guide* of your storage for installation instructions of AIX Host Attachment.

---

## Installing and upgrading SDDPCM

Use these topics to install and upgrade SDDPCM.

### Installing SDDPCM from CD-ROM

SDDPCM is released as an AIX installation image. The SDDPCM install image resides in the `/usr/sys/inst.images/SDDPCM` directory on CD-ROM directory. Because the package does not reside in the `/usr/sys/inst.images` directory, which is the default directory for the AIX install program, you must mount the CD-ROM file system before you can use SMIT to install SDDPCM from the CD-ROM directory.

#### Notes:

1. To mount the CD-ROM and install SDDPCM, you must have root access and AIX system administrator knowledge.
2. The AIX SDDPCM host attachment package (`devices.fcp.disk.ibm.mpio.rte`) must be installed before you install the SDDPCM package (`devices.sddpcm.52.rte`, `devices.sddpcm.53.rte`, or `devices.sddpcm.61.rte` package).

### Creating and mounting the CD-ROM filesystem

To install SDDPCM from the CD-ROM, you must first create and mount the CD-ROM filesystem. Use SMIT to perform the following steps to create and mount the CD-ROM to CD-ROM file system.

**Note:** Throughout this procedure, `/dev/cd0` is used for the compact disc driver address. The driver address can be different in your environment.

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

- Click the default option settings for the other fields to ensure that they are what you need.

```

+-----+
+ Add a CDROM File System +
+ +
+ Type or select values in entry fields. +
+ Press Enter AFTER making all desired changes. +
+ +
+ [Entry Fields] +
+ * DEVICE name cd0 +
+ * MOUNT POINT [/cdmnt] +
+ Mount AUTOMATICALLY at system restart? no +
+ +
+-----+

```

- Press **Enter** to create the CDROM File System.
- When the CDROM File System has been created, press **F10** to exit from smit.
- From your desktop window, enter **smitty mount** and press **Enter**.
- Select **Mount a File System** and press **Enter**. The Mount a File System panel is displayed.
- Select **FILE SYSTEM name** and press **F4**
- Select the CDROM File System that you created and press **Enter**.
- Select **DIRECTORY** on which to mount and press **F4**.
- Select the CDROM File System that you created and press **Enter**.
- Select **TYPE of file system** and press **Enter**.
- Select **cdrfs** as the type of file system and press **Enter**.
- Select **Mount as a REMOVABLE file system?** and press **TAB** to change the entry to *yes*.
- Select **Mount as a READ-ONLY system?** and press **TAB** to change entry to *yes*.
- Click to check the default option settings for the other fields to ensure that they are what you need.

```

+-----+
+ Mount a File System +
+ Type or select values in entry fields. +
+ Press Enter AFTER making all desired changes. +
+ [Entry Fields] +
+ FILE SYSTEM name [/dev/cd0] +
+ DIRECTORY over which to mount [/cdmnt] +
+ TYPE of file system cdrfs +
+ FORCE the mount? no +
+ REMOTE NODE containing the file system [] +
+ to mount +
+ Mount as a REMOVABLE file system? yes +
+ Mount as a READ-ONLY system? yes +
+ Disallow DEVICE access via this mount? no +
+ Disallow execution of SUID and sgid programs no +
+ in this file system? +
+ +
+-----+

```

- Press **Enter** to mount the file system.
- When the file system has been mounted successfully, press **F10** to exit from smit.

## Using the System Management Interface Tool facility to install SDDPCM

Use the System Management Interface Tool (SMIT) facility to install SDDPCM. The SMIT facility has two interfaces: nongraphical (enter **smitty** to invoke the nongraphical user interface) and graphical (enter **smit** to invoke the graphical user interface).

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

1. From your desktop window, `cd` to the directory where the CD-ROM file system is mounted, for example `/cdmnt`.
2. Go to the directory `usr/sys/inst.images/SDDPCM`.
3. From your desktop window, enter **smitty install\_update** and press **Enter** to go directly to the installation panels. The Install and Update Software menu is displayed.
4. Highlight **Install Software** and press **Enter**.
5. Enter `.` to indicate the current directory and press **Enter**.
6. Highlight **Software to Install** and press **F4**. The Software to Install panel is displayed.
7. Select the `devices.sddpcm.52.rte`, `devices.sddpcm.53.rte`, or `devices.sddpcm.61.rte` package, based on the OS level.
8. Press **Enter**. The Install and Update from LATEST Available Software panel is displayed with the name of the software that you selected to install.
9. Check the default option settings to ensure that they are what you need.
10. Press **Enter** to install. SMIT responds with the following message:

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

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

## Unmounting the CD-ROM File System

After successfully installing SDDPCM, use the following procedure to unmount CD-ROM file system in order to remove the CD-ROM:

1. Go to the root (`/`) directory.
2. Enter **umount /cdmnt** and press **Enter** to unmount the CD-ROM file system from the `/cdmnt` directory.
3. Enter **rmfs /cdmnt** and press **Enter** to remove the CD-ROM file system.
4. Remove the CD-ROM.

## Installing SDDPCM from downloaded code

Use this procedure to install SDDPCM from downloaded code:

1. Download code from the Web site that is appropriate for your system
2. Move the downloaded files to the `/usr/sys/inst.images` directory
3. Change directory to the `/usr/sys/inst.images` directory.
4. Untar the downloaded file: for example, `tar -xvf devices.sddpcm.53.rte.tar`
5. Update the `.toc` file. For example:

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

This command should reflect the newer SDDPCM code version that will be updated.

6. Continue the installation by following the instructions beginning in step 3 on page 100.

## Installing SDDPCM with the AIX OS from an AIX NIM SPOT server to the client SAN boot disk or the internal boot disk

You can install SDDPCM from an AIX Network Installation Management (NIM) server to the client SAN boot disk or the internal boot disk at the same time that the AIX OS is installed. You must set up the NIM master and create the lpp\_source and Shared Product Object Tree (SPOT) resources with the images on a file system, which is either NFS-exported or is obtained from a CD or DVD.

Prepare for the NIM SPOT installation with AIX OS and SDDPCM on the client's SAN boot disk or the internal boot disk. To do this, first set up a NIM master and create the lpp\_source and SPOT resource. You can use the System Management Interface Tool (SMIT) facility to implement the following procedures:

1. Install the following filesets to set up the system as an NIM master:

```
bos.sysmgt.min.master
bos.sysmgt.nim.spot
```

2. Initialize the NIM master system by running the **smitty nim\_config\_env** command.
3. Create a new lpp\_source and SPOT resource by running the **smitty nim\_config\_env** command.
4. Add the SDDPCM fileset to the newly created lpp\_source by running the **smitty nim\_task\_inst** command.
5. Create a SPOT from the new lpp\_source by running the **smitty nim\_config\_env** command.
6. Define an NIM client by running the **smitty nim** command.

See the NIM task roadmap on the Web for detailed information on how to complete these tasks:

[publib16.boulder.ibm.com/pseries/en\\_US/aixins/insgdrf/nim\\_roadmap.htm#nim\\_roadmap](http://publib16.boulder.ibm.com/pseries/en_US/aixins/insgdrf/nim_roadmap.htm#nim_roadmap)

After you have successfully prepared for the NIM SPOT installation, you are ready to use the SMIT tool to start the NIM installation on the client system:

1. Run the **smitty nim** command.
  - a. Click **Perform NIM Administration Tasks > Manage Network Install Resource Allocation > Manage Machines > Allocate Network Install Resources**.
  - b. Select the hostname of the client that you defined previously.
  - c. Select the lpp\_source and SPOT resources that you created previously, and then press **Enter**.
2. Run the **smitty nim** command again.

- a. Click **Perform NIM Administration Tasks > Manage Machines > Perform Operations on Machines**.
- b. Select the hostname of the client that you selected previously.
- c. Click **bos\_inst**.
- d. Set the **ACCEPT new license agreements** field to **Yes**, and then press **Enter**.

The system automatically reboots after the **smitty nim** task completes. Use the following command to check the SAN boot disk and make sure the boot disk is configured with SDDPCM:

```
lsattr -El hdiskX (SAN boot disk device name)
```

From the output of this command, check the ODM attribute PCM to ensure that the value is *PCM/friend/sddpcm*.

## Updating SDDPCM

The following sections discuss the following methods of updating SDDPCM:

- “Updating SDDPCM packages by installing a newer base package or a program temporary fix”
- “Committing or rejecting a program temporary fix update” on page 103
- “Verifying the currently installed version of SDDPCM” on page 104
- “Maximum number of devices that SDDPCM supports” on page 104

### Updating SDDPCM packages by installing a newer base package or a program temporary fix

SDDPCM allows you to update SDDPCM by installing a newer base package or a program temporary fix (PTF). A PTF file has a file extension of .bff (for example, *devices.sddpcm.52.rte.2.1.0.1.bff*) and can either be applied or committed when it is installed. If the PTF is committed, the update to SDDPCM is permanent; to remove the PTF, you must uninstall SDDPCM. If the PTF is applied, you can choose to commit or to reject the PTF at a later time. If you decide to reject the PTF, you will not need to uninstall SDDPCM from the host system.

Before applying a newer base package or a PTF to your system, you must stop the applications and stop the SDDPCM server daemon. After applying the base package or the PTF, follow the procedure in “Configuring and unconfiguring supported storage MPIO-capable devices” on page 106 to reconfigure the supported storage devices. The SDDPCM server daemon should be automatically started after restarting the system. If it does not start automatically, you need to start the SDDPCM server daemon manually.

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

If the base package or PTF is on a CD-ROM, you need to mount the CD file system, and then **cd** to the directory on the CD that contains the SDDPCM base package or PTF. See “Creating and mounting the CD-ROM filesystem” on page 98 for directions on how to mount the CD file system. Throughout this SMIT procedure, */dev/cd0* is used for the CD drive address. The drive address can be different in your environment.

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

1. Log in as the root user.
2. From your desktop window, enter **smitty install\_update** and press **Enter** to go directly to the installation panels. The Install and Update Software menu is displayed.
3. Select **Install Software** and press **Enter**.
4. Enter **.** to select the current directory as the INPUT Device/Directory for Software panel and press **Enter**. The Install Software panel is displayed.
5. Select **Software to Install** and press **F4**. The Software to Install panel is displayed.
6. Select the base package or the PTF package that you want to install.
7. Press **Enter**. The Install and Update from LATEST Available Software panel is displayed with the name of the software that you selected to install.
8. If you only want to apply the PTF, select **Commit software Updates?** and tab to change the entry to *no*. The default setting is to commit the PTF. If you specify *no* to Commit Software Updates?, ensure that you specify *yes* to Save Replaced Files?.
9. Check the other default option settings to ensure that they are what you need.
10. Press **Enter** to install. SMIT responds with the following message:

```

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

```

11. Press **Enter** to continue. The installation process can take several minutes to complete.
12. When the installation is complete, press **F10** to exit from SMIT.
13. Unmount the CD-ROM file system and remove the compact disc.

### Committing or rejecting a program temporary fix update

Before you reject a PTF update, you need to unconfigure and remove all supported storage devices from your host system. Committing a PTF does not require this extra step. Perform the following steps to commit or reject a PTF update with the SMIT facility. The SMIT facility has two interfaces: nongraphical (enter **smitty** to invoke the nongraphical user interface) and graphical (enter **smit** to invoke the GUI).

1. Log in as the root user.
2. From your desktop window, enter **smitty install** and press **Enter** to go directly to the installation panels. The Software Installation and Maintenance menu is displayed.
3. Select **Software Maintenance and Utilities** and press **Enter**.
4. Select **Commit Applied Software Updates** to commit the PTF or select **Reject Applied Software Updates** to reject the PTF.
5. Press **Enter**. The Commit Applied Software Updates panel is displayed or the Reject Applied Software Updates panel is displayed.
6. Select **Software name** and press **F4**. The software name panel is displayed.
7. Select the Software package that you want to commit or reject.
8. Check the default option settings to ensure that they are what you need.
9. Press **Enter**. SMIT responds with the following message:

```

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

```

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

**Note:** You do not need to restart the system even though the bosboot message may indicate that a restart is necessary.

### Verifying the currently installed version of SDDPCM

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

```
lslpp -l *sddpcm*
```

or

```
pcmpath query version
```

**Note:** Using the **pcmpath** command requires that at least one SDDPCM-supported MPIO device is configured on the system.

### Maximum number of devices that SDDPCM supports

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

In order to support 1200 supported storage device LUNs, system administrators should first determine whether the system has sufficient resources to support a large number of devices. See “Preparing your system to configure more than 600 supported storage devices or to handle a large amount of I/O after queue depth is disabled” on page 30 for more information.

## Migrating SDDPCM

The following sections discuss the methods of migrating SDDPCM with and without SAN boot devices

- “Migrating the supported storage SAN boot device or nonboot volume group from AIX default PCM to SDDPCM”
- “Migrating from SDDPCM to the AIX default PCM or to SDD” on page 106
- “Migrating from SDD with SAN boot devices (on supported storage hdisks) to SDDPCM with multipath SAN boot devices” on page 106

### Migrating the supported storage SAN boot device or nonboot volume group from AIX default PCM to SDDPCM

The default reserve policy of the AIX base PCM is single-path policy, which is scsi-2 reserve. The path selection algorithm is *fail\_over*, which means only one path is opened at a time and that path made scsi-2 reserve to the disk. All I/O is routed to this path. This reserve policy and path selection algorithm can cause problems if



you build a volume group and file system with AIX default PCM and leave the volume groups active and file system mounted before restarting the system after the SDDPCM packages are installed.

After the system starts, you might see some paths in INVALID state. INVALID state means that the path failed to open. You might not see paths in INVALID state if your system is at AIX 5.2 TL10 or later, AIX 5.3 TL07 or later, or if you have IY83717 or IY83847 installed on your system. Instead, you might see a heavier I/O select count on one or more paths. This is because the scsi-2 reserve is not released during the system restart; thus, only the paths previously opened with scsi-2 reserve can be opened and used for I/O after system restart. All the other paths cannot be opened or used for I/O because of a reservation conflict.

If you have supported storage SAN boot devices configured with AIX default PCM and the reserve policy is single\_path (scsi-2 reserve), switching the boot devices from AIX default PCM to SDDPCM might result in this reservation conflict situation. It is recommended that you always run the **relbootrsv** command to release the scsi-2 reserve on SAN boot devices after you install the SDDPCM host attachment package and the SDDPCM package and before you restart the system. After you run the **relbootrsv** command, run the following command against the hdisks that are part of the rootvg to verify that they are no longer reserved.

```
pcmquerypr -Vh /dev/hdisk6
connection type: fscsi0
open dev: /dev/hdisk6

Attempt to read reservation key...

Attempt to read registration keys...
Read Keys parameter
 Generation : 0
 Additional Length: 0
resrvpolicy= no_reserve
Reserve Key provided by current host = none (hex)02bbb003
Not reserved.
```

To prevent this problem from occurring on nonboot volume groups, perform one of the following actions:

- Switch from the AIX default PCM to SDDPCM before making any volume groups and file systems.
- To switch from the AIX default PCM to SDDPCM, you must unmount file systems and vary off the volume group of the AIX default PCM to release the scsi-2 reserve on the volume group before system restart.
- Issue **relbootrsv VGname** to release the scsi-2 reserve on the active, nonboot volume group devices before you restart the system.

To use the **relbootrsv** command to release a non-SAN boot volume group, issue:  
**>relbootrsv VGname**

To use the **relbootrsv** command to release a SAN boot volume group, issue:  
**>relbootrsv**

**relbootrsv** will search for active SAN boot devices and release the scsi-2 reserve of the boot devices.

## Migrating from SDDPCM to the AIX default PCM or to SDD

**Note:** If you have supported storage devices configured with SDDPCM as MPIO boot devices, you need to contact IBM Customer Support for migration from SDDPCM to the AIX default PCM. The following procedure is applicable only for a system without SDDPCM MPIO devices configured as the SAN boot device.

To migrate from SDDPCM to the AIX default PCM or to SDD, you must first unconfigure the devices, stop the SDDPCM server daemon, and then uninstall the SDDPCM package and the SDDPCM host attachment package. See “Removing SDDPCM from an AIX host system” on page 109 for directions on uninstalling SDDPCM. After you uninstall SDDPCM, you can then restart the system to migrate supported storage MPIO devices to the AIX default PCM. If you want to migrate supported storage devices to SDD devices, you must then install the supported storage device host attachment for SDD and the appropriate SDD package for your system. Then restart the system to configure the supported storage devices to SDD vpath devices.

### Migrating from SDD with SAN boot devices (on supported storage hdisks) to SDDPCM with multipath SAN boot devices

If you have supported storage devices configured with SDD and there are SAN boot devices with supported storage hdisk devices, you need to contact IBM Customer Support for migration from SDD to SDDPCM.

### Migrating SDDPCM during an AIX OS upgrade with multipath SAN boot devices (on supported storage hdisks)

SDDPCM provides different packages to match the AIX OS level. If an AIX system is going to be upgraded to a different OS level, you need to install the corresponding SDDPCM package for that OS level.

If you want to upgrade AIX OS and there are SAN boot devices with SDDPCM supported storage hdisk devices, you need to contact IBM Customer Support for migration from SDDPCM during the OS upgrade.

---

## Configuring and unconfiguring supported storage MPIO-capable devices

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

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

### Configuring supported storage MPIO-capable devices

The newly installed supported storage devices must be configured as MPIO-capable devices before you can use them. Use one of the following commands to configure these devices:

- `cfgmgr` command

**Note:** If operating in a switched environment, issue the `cfgmgr` command once for each host adapter each time a device is added.

If you use the **cfgmgr** command to configure supported storage MPIO devices, you might need to start the SDDPCM server daemon manually, if it has not already started. See “SDDPCM server daemon” on page 118 for information describing how to check the daemon status and how to manually start the daemon.

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

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

## Verifying the SDDPCM configuration

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

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

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

1. Enter **smitty MPIO** from your desktop window. The **MPIO management** menu is displayed.
2. Select MPIO Device Management and press **Enter**. The MPIO Device Management panel is displayed.
3. Select List ALL MPIO Devices and press **Enter**. All MPIO devices on the host are listed.
4. Search for all IBM MPIO FC XXXX devices, where XXXX can be 2105, 2107, 1750, or SAN Volume Controller MPIO devices. Ensure that they are in the *Available* state.

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

### Notes:

1. If none of the supported storage devices are configured successfully as MPIO devices, the **pcmpath query device** command will fail.
2. Because the AIX disk-configure method supports parallel device configuration, the same MPIO logical device name can be configured with different supported storage devices after each device configuration. If you have a program or script that has the MPIO logical device name hardcoded, ensure that you check the MPIO logical device name binding with the storage device LUN. If there is a change to the binding between an MPIO logical device name and the storage device LUN, you must revise your program or script.

## Dynamically adding and removing paths or adapters

When supported storage devices are configured as MPIO-capable devices under AIX 5.2 TL07 (or later) or AIX 5.3 TL03 (or later), you can add or remove extra paths or adapters while I/O is running. To add extra paths that are attached to an adapter to existing available devices, enter:

```
mkpath -l hdiskX -p fscsiY
```

When the command returns successfully, the paths are added to the devices.

To check the device configuration status, enter:

```
lspath -l hdiskX -H -F "name path_id parent connection status"
```

or

**pcmpath query device X**

To add a new adapter to existing available supported storage MPIO devices, enter:  
**cfgmgr -vl fscsiX**

To check the adapter configuration status, enter:

**pcmpath query adapter**

or

**pcmpath query device**

To dynamically remove all paths under a parent adapter from a supported storage MPIO device, enter: **rmpath -dl hdiskX -p fscsiY**

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

**rmdev -l fscsiX -R**

To dynamically remove a particular path, run **smit mpio**, or enter one of the following commands on the command line:

**rmpath -l hdiskX -p fscsiY -w connection location code**

or

**rmpath -dl hdiskX -p fscsiY -w connection location code**

Issue the following command to get a particular path connection location code:

**lspath -l hdiskX -H -F "name path\_id parent connection status"**

**Note:** You cannot remove the last path from a supported storage MPIO device. The command will fail if you try to remove the last path from a supported storage MPIO device.

## Unconfiguring supported storage MPIO-capable devices

**Note:** Contact IBM support if your system has SAN boot MPIO devices configured with SDDPCM on supported storage system disks. The SAN boot MPIO devices cannot be unconfigured with following procedures.

To remove all non-SAN supported storage MPIO-capable devices:

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

**rmdev -dl fcsX -R**

---

## Removing SDDPCM from an AIX host system

Before you remove the SDDPCM package from your AIX host system, all supported storage devices must be unconfigured and removed from your host system, or migrated to the AIX default PCM. The SDDPCM server daemon must be stopped.

**Note:** SDDPCM supports MPIO-supported storage devices as SAN boot devices. If your server has a supported storage device configured with SDDPCM as the SAN boot device, contact IBM support for your SAN boot configuration environment.

To remove SDDPCM, you must remove both the supported storage devices host attachment for SDDPCM and the SDDPCM software packages before you reconfigure supported storage devices or restart the system. Otherwise, the devices can be in the *Defined* state and will not be able to be configured as either MPIO or non-MPIO devices.

After all the supported storage devices are removed or migrated to the AIX default PCM and the SDDPCM server daemon (**pcmsrv**) is stopped, perform the following steps to remove the SDDPCM software package:

1. Enter **smitty deinstall** from your desktop window to go directly to the Remove Installed Software panel.
2. Press **F4** in the SOFTWARE name field to bring up a list of packages and press the **F7** key to select the package to uninstall.
3. Press **Tab** in the **PREVIEW Only?** field to toggle between Yes and No. Select **No** to remove the software package from your AIX host system.

**Note:** If you select **Yes**, the process stops at this point and previews what you are removing. The results of your precheck are displayed without removing the software. If the state for any supported storage MPIO device is either *Available* or *Defined*, the process fails.

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

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

6. Press **Enter** to begin the removal process. This might take a few minutes.
7. When the process is complete, the SDDPCM software package and the supported storage device host attachment for SDDPCM are removed from your system.

---

## SDDPCM support for HACMP with Enhanced Concurrent Mode volume groups

Starting from SDDPCM 2.1.2.0, SDDPCM supports HACMP V5.1, V5.3, and V5.4 on an AIX 5.2 TL07 (or later) and AIX 5.3 TL03 (or later) system with both concurrent and nonconcurrent resource groups. SDDPCM 2.2.0.0 supports HACMP V5.3, and V5.4 on an AIX 6.1 system with both concurrent and nonconcurrent resource groups.

This support requires definition of the shared volume groups as Enhanced Concurrent Mode volume groups, in either concurrent or nonconcurrent resource

groups. This means that no reserve needs to be broken in response to a node failure, and hence any requirement on breaking reserves is removed. A special interaction between HACMP and LVM ensures that if the volume group is used in a nonconcurrent resource group, applications are allowed to access it on one node at a time. Only no\_reserve policy is supported in both concurrent and nonconcurrent resource groups.

The Enhanced Concurrent Mode volume groups are sufficient to ensure high availability. However, if system ECM volume groups are in nonconcurrent resource groups, you should configure your SAN using the following guidelines:

- The interaction between HACMP and LVM to ensure that only one node has nonconcurrent access at a time is *advisory locking*. This is in contrast to the *mandatory locking* provided by SCSI reserves. To ensure that production data is not inadvertently modified by nodes that are not in the HACMP cluster, the following should be done:
  1. Use either physical cabling or zoning to ensure that only HACMP nodes have access to the shared LUNs. That is, non-HACMP nodes should be prevented by hardware from accessing the shared LUNs.
  2. Start HACMP on the cluster nodes at boot time. This ensures that HACMP will activate the appropriate access controls on the shared disk before applications have a chance to modify the access controls.
- Configure disk heartbeating to reduce the likelihood of one node considering the other *dead* and attempting to take over the shared disks. (This is known as a *partitioned cluster*, or *split brain syndrome*). If the shared disks consist of multiple enclosures, use one disk in each enclosure as a heartbeat path.

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

<http://www.ibm.com/systems/storage/disk/ess/>  
<http://www.ibm.com/systems/storage/disk/ds6000/>  
<http://www.ibm.com/systems/storage/disk/ds8000/>  
<http://www.ibm.com/systems/storage/software/virtualization/svc/>

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

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

For HACMP up-to-date ARAP information, go to the following Web site:

<http://www14.software.ibm.com/webapp/set2/sas/f/hacmp/download/aix53.html>

---

## MPIO reserve policies

SDDPCM supports four types of MPIO reserve policies. You can select one of the four reserve policies based on their configuration environment or application needs. The supported reserve policies are:

- No Reserve reservation policy
- Exclusive Host Access Single Path reservation policy
- Persistent Reserve Exclusive Host Access reservation policy
- Persistent Reserve Shared Host Access reservation policy



## **No Reserve reservation policy**

If you set MPIO devices with this reserve policy, there is no reserve being made on MPIO devices. A device without reservation can be accessed by any initiators at any time. Input/output can be sent from all the paths of the MPIO device. This is the default reserve policy of SDDPCM.

## **Exclusive Host Access single-path reservation policy**

This is the scsi-2 reservation policy. If you set this reserve policy for MPIO devices, only the fail\_over path selection algorithm can be selected for the devices. With this reservation policy, an MPIO device only has one path being opened, and a scsi-2 reservation is made by this path on the device. Input/output can only be sent through this path. When this path is broken, another path will be opened and scsi-2 reservation will be made by the new path. All input and output will be routed to this path.

## **Persistent Reserve Exclusive Host Access reservation policy**

If you set an MPIO device with this persistent reserve policy, a persistent reservation is made on this device with a persistent reserve (PR) key. Any initiators who register with the same PR key can access this device. Normally, you should pick a unique PR key for a server. Different servers should have different unique PR key. Input and output is routed to all paths of the MPIO device, because all paths of an MPIO device are registered with the same PR key. In a nonconcurrent clustering environment, such as HACMP, this is the reserve policy that you should select.

Current HACMP clustering software supports no\_reserve policy with Enhanced Concurrent Mode volume group. HACMP support for persistent reserve policies for supported storage MPIO devices is not available.

## **Persistent Reserve Shared Host Access reservation policy**

If you set an MPIO device with this persistent reserve policy, a persistent reservation is made on this device with a persistent reserve (PR) key. However, any initiators that implemented persistent registration can access this MPIO device, even if the initiators are registered with different PR keys. In a concurrent clustering environment, such as HACMP, this is the reserve policy that you should select for sharing resources among multiple servers.

Current HACMP clustering software supports no\_reserve policy with Enhanced Concurrent Mode volume group. HACMP support for persistent reserve policies for supported storage MPIO devices is not available.

---

## **SDDPCM ODM attribute settings**

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

- “SDDPCM ODM attribute default settings”
- “Changing device reserve policies” on page 112
- “Changing the path selection algorithm” on page 112
- “Changing SDDPCM path healthcheck mode” on page 113

## **SDDPCM ODM attribute default settings**

SDDPCM has following default attributes settings:



| Attribute                 | Default value |
|---------------------------|---------------|
| device reserve policy     | no_reserve    |
| path selection algorithm  | load balance  |
| healthcheck mode          | nonactive     |
| healthcheck time interval | 60 seconds    |

## Changing device reserve policies

Use the **chdev** command to change the reserve policy for a device. Because **chdev** requires the device to be unconfigured and reconfigured, this is a disruptive operation.

The following reserve policies can be used with any of the supported path selection algorithms (see “Supported SDDPCM features” on page 91):

- no\_reserve
- persistent reserve exclusive host access
- persistent reserve shared host access

When the reserve policy of a device is exclusive host access single path (scsi-2), the only path selection algorithm supported is the fail\_over algorithm. The fail\_over algorithm selects one path at a time for all I/Os. When the active path fails, then an alternative path is selected. The scsi-2 reserve is reissued by this alternative path

To change the device reserve policy to no\_reserve, enter:

```
chdev -l hdiskX -a reserve_policy=no_reserve
```

If you want to change the reserve policy to one of the persistent reserve policies, you must provide a persistent reserve key at the same time that you change the device policy to one of the persistent reserve types. For example, to change the reserve policy to PR\_shared:

```
chdev -l hdiskX -a PR_key_value=0x1234 -a reserve_policy=PR_shared
```

**Note:** Starting from SDDPCM 2.1.0.0, 2 persistent reserve tools are provided to manage supported storage system MPIO devices. See “Persistent reserve command tools” on page 120 for more information.

## Changing the path selection algorithm

Starting with SDDPCM 2.1.0.0, you can use the **pcmpath set device algorithm** command to dynamically change the path selection algorithm. See “pcmpath set device algorithm” on page 143 for information about this command.

You can also use the **chdev** command to change the path selection algorithm of a device. Because **chdev** requires that the device be unconfigured and then reconfigured, this is a disruptive operation.

Use the following command to change the device path selection algorithm to round robin:

```
chdev -l hdiskX -a algorithm=round_robin
```

You can change the `reserve_policy` and `algorithm` for a device with one command. For example, to change the reserve policy to `no_reserve` and the path selection algorithm to `round robin`:

```
chdev -l hdiskX -a reserve_policy=no_reserve -a algorithm=round_robin
```

## Changing SDDPCM path healthcheck mode

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

`hc_mode`

Healthchecking supports the following modes of operations:

- **Enabled** - When this value is selected, the healthcheck command will be sent to paths that are opened with a normal path mode.
- **Failed** - When this value is selected, the healthcheck command will be sent to paths that are in failed state.
- **Nonactive** - When this value is selected, the healthcheck command will be sent to paths that have no active I/O. This includes paths that are opened or in failed state.

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

Starting with SDDPCM 2.1.0.0, the `pcmpath set device hc_mode` command allows you to dynamically change the path healthcheck mode. See “`pcmpath set device hc_mode`” on page 145 for information about this command.

You can also use the `chdev` command to change the device path healthcheck mode. Because `chdev` requires that the device be unconfigured and then reconfigured, this is a disruptive operation. To change the path healthcheck mode to *failed*, issue following command:

```
chdev -l hdiskX -a hc_mode=failed
```

## Changing SDDPCM path healthcheck time interval

The `hc_interval` attribute will determine how often the paths of a device should be health-checked. The `hc_interval` attribute has a range of values from 0 - 3600 seconds. When a value of 0 is selected, the healthcheck function is disabled. Starting from SDDPCM v2.1.2.3, the default value setting is changed from 20 (seconds) to 60 (seconds).

Starting with SDDPCM 2.1.0.0, the `pcmpath set device hc_interval` command allows you to dynamically change the path healthcheck time interval. See “`pcmpath set device hc_interval`” on page 144 for information about this command.

You can also use the `chdev` command to change the device path healthcheck time interval. Because `chdev` requires that the device be unconfigured and then reconfigured, this is a disruptive operation. To disable the path healthcheck interval function, issue the following command:

```
chdev -l hdiskX -a hc_interval=0
```

**Note:** Currently, the SDDPCM healthcheck function only checks the paths being opened. It does not healthcheck any path that is in the close state. The SDDPCM server daemon healthchecks close\_failed paths. If a device healthcheck function is disabled, the SDDPCM server daemon healthchecks the device's failed path that is opened. See Chapter 12 of the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for more information.

From SDDPCM v2.1.2.3 and later, a new feature of SDDPCM server daemon healthcheck function is introduced. The SDDPCM server daemon automatically starts or stops the healthcheck function on a device if you issue one of the following:

- **pcmpath set device m hc\_interval 0** on the fly to disable a device internal healthcheck function
- **pcmpath set device m hc\_interval n** to enable a device internal healthcheck function on the fly

**Note:** The SDDPCM server daemon only healthcheck FAILED paths. It does not implement healthcheck on opened or idle paths as the SDDPCM internal healthcheck does. This is the difference between the SDDPCM internal healthcheck and the SDDPCM server daemon healthcheck.

---

## AIX 5.2 TL07 (or later) or AIX 5.3 TL03 (or later) fibre-channel device driver features

SDDPCM supports the following fibre channel adapter driver features:

- Fast I/O failure of fibre channel devices
- Fibre-channel dynamic device tracking

### Fast I/O failure of fibre-channel devices

The AIX fibre channel adapter driver supports the new `fc_err_recov` attribute in AIX. If enabled, this attribute can reduce the amount of time needed to fail I/O in certain conditions. By default, this attribute has a value of `delayed_fail`, which means that the fast failure is disabled.

To enable fast failure, you must perform the following steps:

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

You need to have the following Fibre Channel device driver APARs applied to your system in order to have all dynamic tracking and fast fail fixes:

| APAR    | Description                                                    |
|---------|----------------------------------------------------------------|
| IY37183 | Dynamic Tracking and Fast I/O Failure of Fibre Channel Devices |
| IY44139 | Dynamic Tracking: PERM I/O errors after Host HBA Cable Pull    |
| IY44142 | Dynamic Tracking: I/O Hang after Host HBA Cable Pull           |
| IY44342 | Dynamic Tracking: ioctl call may fail after N_Port ID Change   |

|         |                                                                    |
|---------|--------------------------------------------------------------------|
| IY44344 | Dynamic Tracking: Back-to-Back Cable Move May Delay Error Recovery |
| IY45368 | Fast Fail/Dynamic Tracking: FC Device Inaccessible after Move      |
| IY46701 | Dynamic Tracking & MPIO: Multiple Cable Swap Cause Path Failure    |

## Fibre channel dynamic device tracking

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

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

### Notes:

1. This 15 seconds window actually includes the time required to bring up the FC link after the cable is reconnected. For example, if it takes 4 seconds to bring up the FC link, the time allowed to leave the cable disconnected is 11 seconds. I/O will fail if the cable is disconnected longer than 11 seconds.
2. Fibre-channel device dynamic tracking does not support the case when a cable is moved from one adapter on the supported storage device to another free, previously unseen adapter on the supported storage device. The reason is that the world wide port name (WWPN) will be different for that adapter, and the supported storage device cannot track that adapter. The WWPN must be a static identifier of a remote port.

AIX has added a new attribute that must be changed to enable dynamic tracking. By default, dynamic tracking is disabled. To enable dynamic tracking, you must do the following:

1. Change all the child devices of fscsiX on your system to the defined state by issuing the **rmdev -l fscsiX -R** command for all adapters on your system.
2. Issue the **chdev -l fscsiX -a dyntrk=yes** command for all adapters on your system.
3. Issue **cfgmgr** to reconfigure all devices back to the available state.

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

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

Installing APAR IY37183 also installs the file:

`/usr/lpp/bos/README.FIBRE-CHANNEL`

This file has more information about the Dynamic Tracking and Fast I/O Failure features.

---

## Multipath SAN boot support

This section describes how to configure supported storage system MPIO devices as the SAN boot device.

### Configuring supported storage system MPIO devices as the SAN boot device

A supported storage MPIO device can be used as the system boot device. To configure the supported storage device boot device with the SDDPCM module:

1. Select one or more supported storage system devices as the boot device.
2. Install one of the following AIX operating systems on the selected supported storage devices.
  - If the selected supported storage device is ESS, the required operating system is AIX 5.2 TL06 ( or later) or AIX 5.3 TL02 (or later).
  - If the selected supported storage device is DS6000, the required operating system is AIX 5.2 TL07 (or later) or AIX 5.3 TL03 (or later).
  - If the selected supported storage device is DS8000, the required operating system is AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1.
3. Restart the system. The supported storage boot device is configured as an MPIO device with AIX default PCM.
4. Install the supported storage device host attachment for SDDPCM and SDDPCM packages.
5. To release scsi-2 reserve on boot devices, issue the **relbootrsv** command. If you want to release non-rootvg scsi-2 reserves, provide the volume group name as a parameter. For example: **relbootrsv vgname**
6. Restart the system.

All supported storage MPIO devices, including supported storage MPIO SAN boot devices, are now configured with SDDPCM.

When you convert a boot device from the AIX default PCM to SDDPCM, you must issue the **relbootrsv** command, as shown in step 5. If you fail to do so, you might encounter a problem where either all paths of the boot device cannot be opened successfully or they can be opened but cannot be used for I/O. This problem occurs because the AIX default PCM has a default reserve policy of single-path (scsi-2). See “Migrating the supported storage SAN boot device or nonboot volume group from AIX default PCM to SDDPCM” on page 104 for information about solving this problem.

There is a known problem during the SAN boot configuration. After the system is restarted followed by the operating system installation on the supported storage MPIO devices, you might see that some paths of the rootvg are in Failed path state. This can happen even if the system is successfully restarted.

This problem is corrected in AIX 5.2 TL08 or later and AIX 5.3 TL04 or later. Apply the following APARs on these OS levels after the first reboot followed by the operating system installation:

- AIX 5.2 TL08 or later: apply APAR IY83717
- AIX 5.3 TL04 or later: apply APAR IY83847

No APAR is available to correct this problem on AIX52 TL07 and AIX53 TL03. If you configure a SAN boot device with supported storage MPIO devices on one of

these operating system levels and experience this problem, you can manually recover these failed paths by issuing one of the following commands:

- `chpath -s E -l hdiskX -p fscsiX`
- `pcmpath set device M path N online`

---

## Support system dump device with the supported storage system MPIO device

You can choose a supported storage MPIO device to configure with the system primary and secondary dump devices. You can configure the system dump device with the supported storage device boot device, or with the nonboot device. The path selection algorithm for the system dump device will automatically default to *failover\_only* when the system dump starts.

During the system dump, only one path is selected for dump requests. If the first path fails, I/O is routed to the next path being selected.

You must apply all the APARs for AIX 5.2 TL08 and later, or AIX 5.3 TL04 and later.

---

## Dynamically enabling and disabling paths or adapters

This section describes how to dynamically enable and disable paths and adapters.

### Dynamically enabling or disabling a path

There are three ways to dynamically enable (place online) or disable (place offline) a path:

1. Use the following `pcmpath` commands to change the path state:

`pcmpath set device M path N online`

or

`pcmpath set device M path N offline`

2. Use the path control commands provided by AIX.

AIX 5.2 TL04 (or later) or AIX 5.3 TL01 (or later) provides several new path control commands. These commands can be used to configure or remove paths, change path state (enable or disable), and display the path current state.

Use the following AIX path command to change the path state:

`chpath -l hdiskX -s E | D -p fscsiX -w "5005076300c99b0a,5200000000000000"`

**Note:** To find the parent adapter and connect location of a path, use the following AIX path command:

`lspath -l hdiskX -H -F "name path_id parent connection"`

3. Use the smitty MPIO management submenu.
  - a. Enter `smitty MPIO` and press **Enter**. This displays the MPIO Management panel.
  - b. Select MPIO Path Management and press **Enter**. This displays the MPIO Path Management panel.
  - c. Select Enable Paths or Disable paths to enable or disable paths.

## Dynamically enabling or disabling an adapter

The SDDPCM `pcmpath` command can be used to enable (place online) or disable (place offline) an adapter.

To disable an adapter, use the following command:

```
pcmpath set adapter N offline
```

**Note:** SDDPCM reserves the last path of a device. This command will fail if there is any device using the last path attached to this adapter

---

## Using the SDDPCM trace function

SDDPCM supports AIX trace function. The SDDPCM trace ID is 5A7. Trace ID 5A7 traces routine entry, exit, and error paths of the algorithm. To use it, manually turn on the trace function before the program starts to run, then turn off the trace function either after the program stops, or any time you need to read the trace report.

To start the trace function, enter:

```
trace -a -j 5A7
```

To stop the trace function, enter:

```
trcstop
```

To read the report, enter:

```
trcrpt | pg
```

To save the trace data to a file, enter:

```
trcrpt > filename
```

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

---

## SDDPCM server daemon

The SDDPCM server (also referred to as *pcmsrv*) is an integrated component of SDDPCM 2.1.0.0 (or later). This component consists of a UNIX application daemon that is installed in addition to the SDDPCM path control module. The SDDPCM server daemon provides path-recovery function for SDDPCM devices and the First Time Data Capture function. For detailed information about the SDDPCM server daemon, see Chapter 12 of the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

## Verifying if the SDDPCM server has started

After you have installed SDDPCM and restarted the system, verify if the SDDPCM server (`pcmsrv`) has automatically started by entering `lssrc -s pcmsrv`. If `pcmsrv` has automatically started, the output from the `lssrc -s pcmsrv` command looks like this:



| Subsystem | GROUP | PID | Status |
|-----------|-------|-----|--------|
| pcmsrv    |       | NNN | Active |

where *NNN* is the process ID number.

The status of **pcmsrv** should be Active if the SDDPCM server has automatically started. If the SDDPCM server has not started, the status will be *Inoperative*. Go to “Starting the SDDPCM server manually” to proceed.

Because **pcmsrv** is bound to the SDDPCM kernel extension module, **pcmsrv** can fail to start if the SDDPCM is installed and the supported storage MPIO devices have not been configured yet. In this case, you can either restart the system or you can start **pcmsrv** manually after supported storage MPIO devices are configured.

Because **pcmsrv** is bound to the SDDPCM kernel extension module, in order to uninstall or upgrade SDDPCM, you must stop **pcmsrv** so that the SDDPCM kernel extension can be unloaded from the system. During an upgrade, the new SDDPCM kernel extension can be loaded into the system when supported storage MPIO devices are configured.

## Starting the SDDPCM server manually

If **pcmsrv** did not start automatically after you performed the SDDPCM Installation and configured supported storage devices, you can start **pcmsrv** by entering:

```
startsrc -s pcmsrv.
```

Go to “Verifying if the SDDPCM server has started” on page 118 to see if you successfully started the SDDPCM server.

## Stopping the SDDPCM server

You can temporarily disable **pcmsrv** by entering:

```
stopsrc -s pcmsrv.
```

This will stop the current version of **pcmsrv**, but **pcmsrv** will start again if the system is restarted.

The SDDPCM server (**pcmsrv**) recovers failed paths for devices that have disabled the SDDPCM healthcheck function. Furthermore, SDDPCM server provides the First Time Data Capture function, which saves SDDPCM internal trace into a log file. You should not stop **pcmsrv** permanently. However, in certain cases when you want to permanently (not to start even after a system restart) disable **pcmsrv**, you must comment out the following line in the system init table (/etc/inittab):

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

See Chapter 12 of the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for information about the following tasks:

- Enabling or disabling a TCP/IP port for SDDPCM server (“Enabling or disabling the sddsrv or pcmsrv TCP/IP port”)
- Changing to a different port number for the SDDPCM server (“Changing the sddsrv or pcmsrv TCP/IP port number”)

---

## SDDPCM utility programs

This section describes the persistent reserve command tools and the SDDPCM pcmpath commands.

### Persistent reserve command tools

Starting with SDDPCM 2.1.0.0, SDDPCM supports two persistent reserve command tools. The following sections describe the persistent reserve command tools.

#### pcmquerypr

The **pcmquerypr** command provides a set of persistent reserve functions. This command supports the following persistent reserve service actions:

- Read persistent reservation key
- Release persistent reserve
- Preempt-abort persistent reserve
- Clear persistent reserve and registration keys

This command can be issued to all system MPIO devices, including MPIO devices not supported by SDDPCM.

The **pcmquerypr** command can be used in the following situation: the reserve policy of the SDDPCM MPIO devices is set to either persistent reserve exclusive host access (PR\_exclusive) or persistent reserve shared host access (PR\_shared), and persistent reserve has been left on the device by a node, blocking access by another node. The **pcmquerypr** command can be used in this situation to query, preempt, or clear the persistent reserve left by a node or server on the devices.

There are more cases when you might need this tool to solve persistent reserve related problems, such as unexpected persistent reserve left on devices because of failing to release the persistent reserve. Caution must be taken with the command, especially when implementing preempt-abort or clear persistent reserve service action. With preempt-abort service action not only the current persistent reserve key is preempted; it also aborts tasks on the LUN that originated from the initiators that are registered with the preempted key. With clear service action, both persistent reservation and reservation key registrations are cleared from the device.

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

#### Description

The **pcmquerypr** command implements certain SCSI-3 persistent reservation commands on a device. The device can be a supported storage MPIO device. This command supports persistent reserve IN and OUT service actions, such as read, reservation key, release persistent reservation, preempt-abort persistent reservation, or clear persistent reservation and reservation key registrations.

#### Syntax

```
►► pcmquerypr [-p] [-c] [-r] [-v] [-V] [-h /dev/PVname] ◀◀
```

Flags:

- p** If the persistent reservation key on the device is different from the current host reservation key, the existing persistent reservation key on the device is preempted. This option can be issued only when the device is not already open.
- c** If there is a persistent reservation on the device, the persistent reservation is removed and all reservation key registrations on the device are cleared. This option can be issued only when the device is not already open.
- r** Removes the persistent reservation key on the device made by this host. This option can be issued only when the device is not already open.
- v** Displays the persistent reservation key if it exists on the device.
- V** Verbose mode. Prints detailed message.

#### **Return code**

If the command issued without options of `-p`, `-r` or `-c`, the command will return:

- 0** There is no persistent reservation key on the device, or the device is reserved by the current host
- 1** The persistent reservation key is different from the host reservation key
- 2** The command failed.

If the command issued with one of the options of `-p`, `-r` or `-c`, it returns:

- 0** The command was successful.
- 2** The command failed.

#### **Examples**

1. To query the persistent reservation on a device, enter **`pcmquerypr -h /dev/hdisk30`**.  
This command queries the persistent reservation on the device without displaying. If there is a persistent reserve on a disk, it returns 0 if the device is reserved by the current host. It returns 1 if the device is reserved by another host.
2. To query and display the persistent reservation on a device, enter **`pcmquerypr -vh /dev/hdisk30`**.  
Same as Example 1. In addition, it displays the persistent reservation key.
3. To release the persistent reservation if the device is reserved by the current host, enter **`pcmquerypr -rh /dev/hdisk30`**.  
This command releases the persistent reserve if the device is reserved by the current host. It returns 0 if the command succeeds or the device is not reserved. It returns 2 if the command fails.
4. To reset any persistent reserve and clear all reservation key registrations, enter **`pcmquerypr -ch /dev/hdisk30`**.  
This command resets any persistent reserve and clears all reservation key registrations on a device. It returns 0 if the command succeeds, or 2 if the command fails.
5. To remove the persistent reservation if the device is reserved by another host, enter **`pcmquerypr -ph /dev/hdisk30`**.


This command removes an existing registration and persistent reserve from another host. It returns 0 if the command succeeds or if the device is not persistent reserved. It returns 2 if the command fails.

## pcmggenprkey

### Description

The **pcmggenprkey** command can be used to set or clear the PR\_key\_value ODM attribute for all SDDPCM MPIO devices. It also can be used to query and display the reservation policy of all SDDPCM MPIO devices and the persistent reserve key, if those devices have a PR key.

### Syntax



### Examples

1. To set the persistent reserve key to all SDDPCM MPIO devices with a provided key value, issue **pcmggenprkey -u -k 0x1234567890abcdef**. This creates a customized PR\_key\_value attribute with the provided key value for all SDDPCM MPIO devices, except the devices that already have the same customized PR key attribute. The provided key must contain either a decimal integer or a hexadecimal integer.
2. To clear the PR\_key\_value attribute from all SDDPCM MPIO devices, issue **pcmggenprkey -u -k none**.
3. To update the customized PR\_key\_value attribute with the HACMP-provided Preserve key or the output string from the **uname** command for all the SDDPCM MPIO devices, issue **pcmggenprkey -u**. When the **-u** option is used without the **-k** option, this command searches for the HACMP-provided Preservekey attribute and uses that value as the PR key if that attribute is available; otherwise, it uses the output string from the **uname** command as the PR key.
4. To display the reserve\_policy, the PR\_key\_value attribute, and the persistent reserve key attribute of all the SDDPCM devices, issue **pcmggenprkey -v**. If the MPIO device does not have a persistent reserve key, a value of *none* is displayed.

## Using SDDPCM pcmppath commands

SDDPCM provides commands that you can use to display the status of adapters that are used to access managed devices, to display the status of devices that the device driver manages, or to map supported storage MPIO devices or paths to a supported storage device location. You can also set individual path conditions either to online or offline, set all paths that are connected to an adapter either to online or offline, or set all paths that are connected to a supported storage device port or ports to online or offline. This section includes descriptions of these commands. Table 12 on page 123 provides an alphabetical list of these commands, a brief description, and where to go in this chapter for more information.

Table 12. Commands

| Command                        | Description                                                                                | Page |
|--------------------------------|--------------------------------------------------------------------------------------------|------|
| pcmpath clear device count     | Dynamically clears the error count or error/select counts to zero.                         | 124  |
| pcmpath disable ports          | Places paths connected to certain ports offline.                                           | 125  |
| pcmpath enable ports           | Places paths connected to certain ports online.                                            | 125  |
| pcmpath open device path       | Opens an INVALID path.                                                                     | 129  |
| pcmpath query adapter          | Displays information about adapters.                                                       | 131  |
| pcmpath query adaptstats       | Displays performance information for all FCS adapters that are attached to SDDPCM devices. | 132  |
| pcmpath query device           | Displays information about devices.                                                        | 133  |
| pcmpath query devstats         | Displays performance information for a single SDDPCM device or all SDDPCM devices.         | 136  |
| pcmpath query essmap           | Displays each device, path, location, and attributes.                                      | 138  |
| pcmpath query portmap          | Displays supported storage MPIO device port location.                                      | 139  |
| pcmpath query version          | Displays the version of the currently installed SDDPCM.                                    | 140  |
| pcmpath query wwpn             | Displays the world wide port name (WWPN) for all fibre-channel adapters.                   | 141  |
| pcmpath set adapter            | Sets all device paths that are attached to an adapter to online or offline.                | 142  |
| pcmpath set device path        | Sets the path of a device to online or offline.                                            | 146  |
| pcmpath set device algorithm   | Set all or some of supported storage MPIO device path selection algorithm                  | 143  |
| pcmpath set device hc_interval | Set all or some of supported storage MPIO device health check time interval                | 144  |
| pcmpath set device hc_mode     | Set all or some of supported storage MPIO device health check mode                         | 145  |

## pcmpath clear device count

The **pcmpath clear device count** command dynamically sets the error counter or error/select counters to zero.

### Syntax:

```
▶▶ pcmpath clear device number 1 count error
 └── device number 2 ─┘ └── all ─┘
```

### Parameters:

*device number 1* <*device number 2*>

When two device numbers are entered, this command applies to all the devices whose index numbers fit within the range of these two device index numbers.

*error*

Clears the error counter of the specified SDDPCM MPIO device or devices.

*all* Clears both the select counter and the error counter of the specified SDDPCM MPIO device or devices.

**Examples:** If you have a non-zero select counter or error counter, entering **pcmpath query device 20** causes the following output to be displayed:

```
DEV#: 20 DEVICE NAME: hdisk20 TYPE: 2145 ALGORITHM: Load Balance
SERIAL: 60050768018180235800000000000463
=====
```

| Path# | Adapter/Path Name | State | Mode   | Select | Errors |
|-------|-------------------|-------|--------|--------|--------|
| 0     | fscsi1/path0      | CLOSE | NORMAL | 14     | 0      |
| 1*    | fscsi1/path1      | CLOSE | NORMAL | 8      | 0      |
| 2     | fscsi3/path2      | CLOSE | NORMAL | 10009  | 0      |
| 3*    | fscsi3/path3      | CLOSE | NORMAL | 8      | 0      |

If you enter the **pcmpath clear device 20 count all** and then enter **pcmpath query device 20**, the following output is displayed:

```
DEV#: 20 DEVICE NAME: hdisk20 TYPE: 2145 ALGORITHM: Load Balance
SERIAL: 60050768018180235800000000000463
=====
```

| Path# | Adapter/Path Name | State | Mode   | Select | Errors |
|-------|-------------------|-------|--------|--------|--------|
| 0     | fscsi1/path0      | CLOSE | NORMAL | 0      | 0      |
| 1*    | fscsi1/path1      | CLOSE | NORMAL | 0      | 0      |
| 2     | fscsi3/path2      | CLOSE | NORMAL | 0      | 0      |
| 3*    | fscsi3/path3      | CLOSE | NORMAL | 0      | 0      |

## pcmpath disable ports

The **pcmpath disable ports** command sets SDDPCM MPIO device paths offline for the specified supported storage system location code.

### Syntax:

```
▶▶—pcmpath disable ports—connection-ess essid—▶▶
```

### Parameters:

#### *connection*

The connection code must be in one of the following formats:

- Single port = R1-Bx-Hy-Zz
- All ports on card = R1-Bx-Hy
- All ports on bay = R1-Bx

Use the output of the **pcmpath query essmap** command to determine the connection code.

#### *essid*

The supported storage device serial number, given by the output of **pcmpath query portmap** command.

**Examples:** If you enter the **pcmpath disable ports R1-B1-H3 ess 12028** command and then enter the **pcmpath query device** command, the following output is displayed:

```
DEV#: 3 DEVICE NAME: hdisk3 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20712028
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi0/path0 CLOSE OFFLINE 6 0
 1 fscsi0/path1 CLOSE NORMAL 9 0
 2 fscsi1/path2 CLOSE OFFLINE 11 0
 3 fscsi1/path3 CLOSE NORMAL 9 0

DEV#: 4 DEVICE NAME: hdisk4 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20712028
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi0/path0 CLOSE OFFLINE 8702 0
 1 fscsi0/path1 CLOSE NORMAL 8800 0
 2 fscsi1/path2 CLOSE OFFLINE 8816 0
 3 fscsi1/path3 CLOSE NORMAL 8644 0

DEV#: 5 DEVICE NAME: hdisk5 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20912028
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi0/path0 CLOSE OFFLINE 8917 0
 1 fscsi0/path1 CLOSE NORMAL 8919 0
 2 fscsi1/path2 CLOSE OFFLINE 9008 0
 3 fscsi1/path3 CLOSE NORMAL 8944 0

DEV#: 6 DEVICE NAME: hdisk6 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20B12028
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi0/path0 CLOSE OFFLINE 9044 0
 1 fscsi0/path1 CLOSE NORMAL 9084 0
 2 fscsi1/path2 CLOSE OFFLINE 9048 0
 3 fscsi1/path3 CLOSE NORMAL 8851 0
```



DEV#: 7 DEVICE NAME: hdisk7 TYPE: 2105E20 ALGORITHM: Load Balance  
SERIAL: 20F12028

```
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi0/path0 CLOSE OFFLINE 9089 0
 1 fscsi0/path1 CLOSE NORMAL 9238 0
 2 fscsil/path2 CLOSE OFFLINE 9132 0
 3 fscsil/path3 CLOSE NORMAL 9294 0
```

DEV#: 8 DEVICE NAME: hdisk8 TYPE: 2105E20 ALGORITHM: Load Balance  
SERIAL: 21012028

```
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi0/path0 CLOSE OFFLINE 9059 0
 1 fscsi0/path1 CLOSE NORMAL 9121 0
 2 fscsil/path2 CLOSE OFFLINE 9143 0
 3 fscsil/path3 CLOSE NORMAL 9073 0
```

## pcmpath enable ports

The **pcmpath enable ports** command sets MPIIO device paths online for the specified supported storage device location code.

### Syntax:

```
▶▶—pcmpath enable ports—connection—ess essid—▶▶
```

### Parameters:

#### *connection*

The connection code must be in one of the following formats:

- Single port = R1-Bx-Hy-Zz
- All ports on card = R1-Bx-Hy
- All ports on bay = R1-Bx

Use the output of the **pcmpath query essmap** command to determine the connection code.

#### *essid*

The supported storage device serial number, given by the output of **pcmpath query portmap** command.

**Examples:** If you enter the **pcmpath enable ports R1-B1-H3 ess 12028** command and then enter the **pcmpath query device** command, the following output is displayed:

```
DEV#: 3 DEVICE NAME: hdisk3 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20112028
```

```
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi0/path0 CLOSE NORMAL 6 0
 1 fscsi0/path1 CLOSE NORMAL 9 0
 2 fscsil/path2 CLOSE NORMAL 11 0
 3 fscsil/path3 CLOSE NORMAL 9 0
```

```
DEV#: 4 DEVICE NAME: hdisk4 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20712028
```

```
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi0/path0 CLOSE NORMAL 8702 0
 1 fscsi0/path1 CLOSE NORMAL 8800 0
 2 fscsil/path2 CLOSE NORMAL 8816 0
 3 fscsil/path3 CLOSE NORMAL 8644 0
```

```
DEV#: 5 DEVICE NAME: hdisk5 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20912028
```

```
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi0/path0 CLOSE NORMAL 8917 0
 1 fscsi0/path1 CLOSE NORMAL 8919 0
 2 fscsil/path2 CLOSE NORMAL 9008 0
 3 fscsil/path3 CLOSE NORMAL 8944 0
```

```
DEV#: 6 DEVICE NAME: hdisk6 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20B12028
```

```
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi0/path0 CLOSE NORMAL 9044 0
 1 fscsi0/path1 CLOSE NORMAL 9084 0
 2 fscsil/path2 CLOSE NORMAL 9048 0
 3 fscsil/path3 CLOSE NORMAL 8851 0
```

DEV#: 7 DEVICE NAME: hdisk7 TYPE: 2105E20 ALGORITHM: Load Balance  
SERIAL: 20F12028

```
=====
```

| Path# | Adapter/Path Name | State | Mode   | Select | Errors |
|-------|-------------------|-------|--------|--------|--------|
| 0     | fscsi0/path0      | CLOSE | NORMAL | 9089   | 0      |
| 1     | fscsi0/path1      | CLOSE | NORMAL | 9238   | 0      |
| 2     | fscsil/path2      | CLOSE | NORMAL | 9132   | 0      |
| 3     | fscsil/path3      | CLOSE | NORMAL | 9294   | 0      |

DEV#: 8 DEVICE NAME: hdisk8 TYPE: 2105E20 ALGORITHM: Load Balance  
SERIAL: 21012028

```
=====
```

| Path# | Adapter/Path Name | State | Mode   | Select | Errors |
|-------|-------------------|-------|--------|--------|--------|
| 0     | fscsi0/path0      | CLOSE | NORMAL | 9059   | 0      |
| 1     | fscsi0/path1      | CLOSE | NORMAL | 9121   | 0      |
| 2     | fscsil/path2      | CLOSE | NORMAL | 9143   | 0      |
| 3     | fscsil/path3      | CLOSE | NORMAL | 9073   | 0      |

## pcmpath open device path

The **pcmpath open device path** command dynamically opens a path that is in Invalid state. You can use this command to open an Invalid path even when I/O is actively running on the devices.

### Syntax:

```
▶▶—pcmpath open device—device number—path—path number————▶▶
```

### Parameters:

#### *device number*

The logical device number of this hdisk, as displayed by the **pcmpath query device** command.

#### *path number*

The path ID that you want to change, as displayed under “Path Name” by the **pcmpath query device** command.

**Examples:** If you enter the **pcmpath query device 23** command, the following output is displayed:

```
DEV#: 23 DEVICE NAME: hdisk23 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20112028
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi1/path0 OPEN NORMAL 557 0
 1 fscsi1/path1 OPEN NORMAL 568 0
 2 fscsi0/path2 INVALID NORMAL 0 0
 3 fscsi0/path3 INVALID NORMAL 0 0
```

Note that the current state of path 2 and path 3 are INVALID, which means that open path 2 and path 3 failed.

If the root cause of the path 2 open failure is fixed and you enter the **pcmpath open device 23 path 2** command, the following output is displayed:

```
Success: device 23 path 2 opened
```

```
DEV#: 23 DEVICE NAME: hdisk23 TYPE: 2105E20 ALGORITHM: Load Balance
SERIAL: 20112028
=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi1/path0 OPEN NORMAL 557 0
 1 fscsi1/path1 OPEN NORMAL 568 0
 2 fscsi0/path2 OPEN NORMAL 0 0
 3 fscsi0/path3 INVALID NORMAL 0 0
```

After issuing the **pcmpath open device 23 path 2** command, the state of path 2 becomes OPEN.

The terms used in the output are defined as follows:

**Dev#** The logical device number of this hdisk.

**Device name**  
The name of this device.

**Type** The device product ID from inquiry data.

**Algorithm**

The current path selection algorithm for the device. The algorithm selected is one of the following: load balancing, round robin, or failover.

**Serial** The LUN for this device.

**Path#** The path index displayed by the **pcmpath query device** command.

**Adapter**

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

**Path Name**

The name of the path. The number displayed as part of the name is the path ID of this path that is used by the **pcmpath open device path** and **pcmpath set device path** commands.

**State** The condition of each path of the named device:

**Open** Path is in use.

**Close** Path is not being used.

**Close\_Failed** Path is broken and is not being used.

**Failed** Path is opened, but no longer functional because of error.

**Invalid** The path failed to open.

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

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

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

## pcmpath query adapter

The **pcmpath query adapter** command displays information about a single adapter or all adapters that are attached to SDDPCM-configured MPIO devices.

### Syntax:

►►—pcmpath query adapter—*adapter number*—————►►

### Parameters:

*adapter number*

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

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

**Active Adapters :2**

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

The terms used in the output are defined as follows:

#### **Adpt #**

The index number of the adapter.

**Name** The name of the adapter.

**State** The condition of the named adapter. It can be either:

**Normal** Adapter is in use.

**Degraded** One or more opened paths are not functioning.

**Failed** All opened paths that are attached to this adapter are not functioning.

**Mode** The mode of the named adapter, which is either *Active* or *Offline*.

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

**Errors** The number of errors that occurred on all paths that are attached to this adapter.

**Paths** The number of paths that are attached to this adapter.

**Active** The number of functional paths that are attached to this adapter. The number of functional paths is equal to the number of opened paths attached to this adapter minus any that are identified as failed or disabled (offline).

## pcmpath query adaptstats

The **pcmpath query adaptstats** command displays information about a single or all fibre-channel adapters that are attached to SDDPCM-configured MPIO devices. If you do not enter a device number, information about all devices is displayed.

### Syntax:

```
►►—pcmpath query adaptstats—adapter number—————►►
```

### Parameters:

*adapter number*

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

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

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

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

The terms used in the output are defined as follows:

#### Total Read

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

#### Total Write

- I/O: total number of completed Write requests
- SECTOR: total number of sectors that have been written

#### Active Read

- I/O: total number of Read requests in process
- SECTOR: total number of sectors to read in process

#### Active Write

- I/O: total number of Write requests in process
- SECTOR: total number of sectors to write in process

#### Maximum

- I/O: the maximum number of queued I/O requests
- SECTOR: the maximum number of queued sectors to Read or Write



## pcmpath query device

The **pcmpath query device** command displays information about a single SDDPCM MPIO device, a set of SDDPCM MPIO devices between device M and N, a specific model of SDDPCM MPIO devices, or all SDDPCM MPIO devices. If two device numbers are entered, the command displays device information for a set of devices between these two device numbers. The option to specify a model cannot be used when you query with device numbers. If you do not enter a device number, information about all devices is displayed. If a device number is entered, the command displays the device information about the hdisk that is associated with this number.

Starting from SDDPCM 2.1.0.7, a new option, **-d**, is provided to display devices of a particular device model. The valid device models are:

- 2105 - Display all 2105 models (ESS)
- 2107 - Display all 2107 models (DS8000)
- 1750 - Display all 1750 models (DS6000)
- 2145 - Display all 2145 models (SAN Volume Controller)

Starting with SDDPCM 2.1.0.7, a new feature is available to verify SAN configurations in a controller environment, such as DS6000 or SAN Volume Controller. A nonpreferred path will be marked with **\*\***.

For example, if you configure 4 paths per MPIO hdisk device and you want to configure an equal distribution between preferred controller and nonpreferred controller of DS6000 devices, you can configure your environment with 2 paths connected to a preferred controller and 2 paths connected to a nonpreferred controller. This feature will help you verify your configuration by indicating which path is connected to the nonpreferred controller before the device starts operation.

**Note:** You must open SDDPCM MPIO devices once before the **\*\*** mark on nonpreferred paths can be displayed in following 3 commands:

- **pcmpath query device**
- **pcmpath query essmap**
- **pcmpath query portmap**

Starting from SDDPCM 2.1.3.0, two new options are added to the **device query** command. The first option lets you specify two numbers to query a set of devices; the second option **-i x y** lets you repeat the query command every **x** seconds for **y** times.

The **pcmpath query device** commands display only supported storage MPIO devices that are configured with the SDDPCM module. Any AIX internal disks or non-SDDPCM-configured MPIO devices are not displayed.

### Syntax:

```
▶▶—pcmpath query device device number
| device number m device number n
| -d device model
| -i x
| -i x y
```

### Parameters:

*device number*

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

*device number\_m device\_number\_n*

Use the *device\_number\_m device\_number\_n* option to provide a range of device index numbers.

*device model*

Displays devices of a particular device model. The valid device models are:

- 2105 - Display all 2105 models (ESS)
- 2107 - Display all 2107 models (DS8000)
- 1750 - Display all 1750 models (DS6000)
- 2145 - Display all 2145 models (SAN Volume Controller)

**i** Repeats the command every *x* seconds for *y* times. If you do not specify *y*, the command repeats indefinitely every *x* seconds.

**Examples:** If you enter the **pcmpath query device 10** command, the following output is displayed:

For the supported storage device:

DEV#: 10 DEVICE NAME: hdisk10 TYPE: 2105800 ALGORITHM:Load Balance  
SERIAL: 7BFFCA30

```

=====
Path# Adapter/Path Name State Mode Select Errors
 0 fscsi1/path2 CLOSE NORMAL 0 0
 1 fscsi1/path3 CLOSE NORMAL 0 0
 2 fscsi0/path0 CLOSE NORMAL 0 0
 3 fscsi0/path1 CLOSE NORMAL 0 0
=====

```

The terms used in the output are defined as follows:

**Dev#** The logical device number of this hdisk.

**Name** The logical name of this device.

**Type** The device product ID from inquiry data.

**Algorithm**

The current path selection algorithm selected for the device. The algorithm selected is one of the following: load balancing, round robin, or failover.

**Serial** The LUN for this device.

**Path#** The path index displayed by device query command.

**Adapter**

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

**Path Name**

The name of the path. The number displayed as part of the name is the path ID that is used by **pcmpath open device path** and **pcmpath set device path** commands.

**State** The condition of the path attached to the named device:

**Open** Path is in use.

**Close** Path is not being used.

**Failed** Path is no longer being used. It has been removed from service due to errors.

**Close\_Failed**

Path was detected to be broken and failed to open when the device was opened. The path stays in Close\_Failed state when the device is closed.

**Invalid**

The path is failed to open, but the MPIO device is opened.

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

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

**Errors** The number of input and output errors that occurred on a path of this device.

## pcmpath query devstats

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

**Note:** The **pcmpath query devstats** command displays only SDDPCM MPIO devices. Any AIX internal disks or non-SDDPCM configured MPIO devices will not be displayed.

Starting from SDDPCM 2.1.3.0, two new options are added to the **query devstats** command. The first option lets you specify two numbers to query a set of devices; the second option **-i x y** lets you repeat the query command every *x* seconds for *y* times.

### Syntax:

```
▶▶—pcmpath query devstats— device number—————▶▶
|— device number m device number n—————|
|— -d device model—————|
|— -i x—————|
|— -i x y—————|
```

### Parameters:

#### *device number*

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

#### *device number\_m device\_number\_n*

Use the *device\_number\_m device\_number\_n* option to provide a range of device index numbers.

#### *device model*

Displays devices of a particular device model. The valid device models are:

- 2105 - Display all 2105 models (ESS)
- 2107 - Display all 2107 models (DS8000)
- 1750 - Display all 1750 models (DS6000)
- 2145 - Display all 2145 models (SAN Volume Controller)

**i** Repeats the command every *x* seconds for *y* times. If you do not specify *y*, the command repeats indefinitely every *x* seconds.

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

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

Transfer Size: <= 512 <= 4k <= 16K <= 64K > 64K
 30 40 0 0 0
/*-----*/
```

The terms used in the output are defined as follows:

**Total Read**

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

**Total Write**

- I/O: total number of completed Write requests
- SECTOR: total number of sectors that have been written

**Active Read**

- I/O: total number of Read requests in process
- SECTOR: total number of sectors to read in process

**Active Write**

- I/O: total number of Write requests in process
- SECTOR: total number of sectors to write in process

**Maximum**

- I/O: the maximum number of queued I/O requests
- SECTOR: the maximum number of queued sectors to Read or Write

**Transfer size**

- <= 512: the number of I/O requests received, whose transfer size is 512 bytes or less
- <= 4k: the number of I/O requests received, whose transfer size is 4 KB or less (where KB equals 1024 bytes)
- <= 16K: the number of I/O requests received, whose transfer size is 16 KB or less (where KB equals 1024 bytes)
- <= 64K: the number of I/O requests received, whose transfer size is 64 KB or less (where KB equals 1024 bytes)
- > 64K: the number of I/O requests received, whose transfer size is greater than 64 KB (where KB equals 1024 bytes)

## pcmpath query essmap

The **pcmpath query essmap** command displays data collected from the **fcppcmmap** program after the system has been restarted, or by the **pcmpath** program if any device configuration change occurs after the system has been restarted. The **fcppcmmap** program issues SCSI commands directly to the disks to collect the information that is displayed.

### Syntax:

►—pcmpath query essmap—◄

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

| Disk   | Path  | P | Location     | adapter | LUN SN      | Type         | Size | LSS | Vol | Rank | C/A | S ... |
|--------|-------|---|--------------|---------|-------------|--------------|------|-----|-----|------|-----|-------|
| hdisk5 | path0 | * | 30-60-01[FC] | fscsi1  | 13AAAAA1200 | IBM 1750-500 | 1.1  | 18  | 0   | 0000 | 01  | Y ... |
| hdisk5 | path1 |   | 30-60-01[FC] | fscsi0  | 13AAAAA1200 | IBM 1750-500 | 1.1  | 18  | 0   | 0000 | 01  | Y ... |
| hdisk5 | path2 | * | 20-60-01[FC] | fscsi0  | 13AAAAA1200 | IBM 1750-500 | 1.1  | 18  | 0   | 0000 | 01  | Y ... |
| hdisk5 | path3 |   | 20-60-01[FC] | fscsi1  | 13AAAAA1200 | IBM 1750-500 | 1.1  | 18  | 0   | 0000 | 01  | Y ... |

The terms used in the output are defined as follows:

|                   |                                                                                                                                     |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>Disk</b>       | The logical device name assigned by the host.                                                                                       |
| <b>Path</b>       | The logical path name of a MPIO device.                                                                                             |
| <b>P</b>          | Indicates the logical paths and whether the path is preferred and nonpreferred. "*" indicates that the path is a nonpreferred path. |
| <b>Location</b>   | The physical location code of the host adapter through which the LUN is accessed.                                                   |
| <b>Adapter</b>    | The logical adapter name assigned by the host LUN.                                                                                  |
| <b>LUN SN</b>     | The unique serial number for each LUN within the supported storage device.                                                          |
| <b>Type</b>       | The device and model.                                                                                                               |
| <b>Size</b>       | The capacity of the configured LUN.                                                                                                 |
| <b>LSS</b>        | The logical subsystem where the LUN resides. (Starting from 2.1.3.0, the value displayed is changed from decimal to hexadecimal.)   |
| <b>Vol</b>        | The volume number within the LSS.                                                                                                   |
| <b>Rank</b>       | The unique identifier for each RAID array within the supported storage device.                                                      |
| <b>C/A</b>        | The cluster and adapter accessing the array.                                                                                        |
| <b>S</b>          | Indicates that the device is shared by two and more supported storage device ports. Valid values are <i>yes</i> or <i>no</i> .      |
| <b>Connection</b> | The physical location code of supported storage device adapter through which the LUN is accessed.                                   |
| <b>Port</b>       | The supported storage device port through which the LUN is accessed.                                                                |
| <b>RaidMode</b>   | The disk RAID mode.                                                                                                                 |

## pcmpath query portmap

The **pcmpath query portmap** command displays hdisk status from a supported storage device physical location code perspective.

### Syntax:

►►—pcmpath query portmap—►►

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

```

 BAY-1(B1) BAY-2(B2) BAY-3(B3) BAY-4(B4)
 ESSID DISK H1 H2 H3 H4 H1 H2 H3 H4 H1 H2 H3 H4 H1 H2 H3 H4
 ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD
 BAY-5(B5) BAY-6(B6) BAY-7(B7) BAY-8(B8)
 H1 H2 H3 H4 H1 H2 H3 H4 H1 H2 H3 H4 H1 H2 H3 H4
 ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD ABCD

 13AAAKA hdisk5 0--- ---- ---- ---- o--- ---- ---- ---- ---- ---- ---- ---- ---- ---- ---- ----
 13AAAKA hdisk6 Y--- ---- ---- ---- y--- ---- ---- ---- ---- ---- ---- ---- ---- ---- ---- ----

 Y = online/open y = (alternate path) online/open
 O = online/closed o = (alternate path) online/closed
 N = offline n = (alternate path) offline
 - = path not configured
 PD = path down

```

The terms used in the output are defined as follows:

- Y** The port is online and open, meaning that at least one path attached to this port is functional.
- y** Paths connected to this port are nonpreferred paths. The port is online and open, meaning that at least one path attached to this port is functional.
- O** The port is online and closed, meaning that at least one path state and mode is closed and online.
- o** Paths connected to this port are nonpreferred paths. The port is online and closed, meaning that at least one path state and mode is closed and online.
- N** The port is offline, meaning that all paths attached to this port are offline.
- n** Paths connected to this port are nonpreferred paths. The port is offline, meaning that all paths attached to this port are offline.
- The path is not configured.
- PD** The path is down. It is either not functional or has been placed offline.

**Note:** The following fields apply only 1750 devices and can only be shown after the device is opened once:

- y
- o
- n

The serial number of ESS devices is five digits, whereas the serial number of DS6000 and DS8000 devices is seven digits.



## **pcmpath query version**

The **pcmpath query version** command displays the version of SDDPCM that is installed.

### **Syntax:**

▶▶—pcmpath query version—————▶▶

**Parameters:** None

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

```
[root@abc]> pcmpath query version
IBM SDDPCM Version 2.1.1.0 (devices.sddpcm.52.rte)
```

## **pcmpath query wwpn**

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

### **Syntax:**

▶▶—pcmpath query wwpn—————▶▶

**Parameters:** None

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

| Adapter Name | PortWWN          |
|--------------|------------------|
| fscsi0       | 10000000C925F5B0 |
| fscsi1       | 10000000C9266FD1 |

## pcmpath set adapter

The **pcmpath set adapter** command sets all device paths attached to an adapter either to online or offline. SDDPCM reserves the last path of a device if the device is opened. This command will fail if there is any device that has the last path attached to this adapter. If all paths of a device are attached to a single fibre-channel adapter that connects to multiple storage ports through a switch, the **pcmpath set adapter 0 offline** command will fail, because the device will lose all its paths.

**Attention:** If the device is closed, the **pcmpath set adapter offline** command does not have the last path offline protection.

### Syntax:

```
▶▶ pcmpath set adapter-adapter number [online | offline] ▶▶
```

### Parameters:

*adapter number*

The index number of the adapter that you want to change.

*online*

Enables the adapter for service.

*offline*

Disables the adapter from service.

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

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

## pcmpath set device algorithm

The **pcmpath set device algorithm** command dynamically changes the path selection algorithm per MPIO device.

### Syntax:

```
▶▶—pcmpath set device— num1—└─┬─┘—algorithm-option—▶▶
 └─┬─┘
 num2
```

**Note:** If device reserve policy is set to `single_path` ( `scsi-2` reserve), the device algorithm must be set to `fail_over`. Any attempt to set algorithm to `round_robin` or `load_balance` with `single_path` reserve policy will fail.

### Parameters:

*num1* [ *num2* ]

- When only *num1* is specified, the command applies to the hdisk specified by *num1*.
- When 2 device logical numbers are entered, this command applies to all the devices whose logical numbers fit within the range of the two device logical numbers.

*option*

Specifies one of the following path selection algorithms:

- **rr**, where *rr* indicates round robin
- **lb**, where *lb* indicates load balancing
- **fo**, where *fo* indicates failover policy

**Note:** You can enter the **pcmpath set device *N* algorithm *rr/fo/lb*** command to dynamically change to the path selection algorithm associated with SDDPCM MPIO devices that are in either Close or Open state.

**Examples:** If you enter **pcmpath set device 2 10 algorithm rr**, the path-selection algorithm of hdisk 2 to hdisk 10 is immediately changed to the round robin algorithm.

## pcmpath set device hc\_interval

The **pcmpath set device hc\_interval** command dynamically changes the health check time interval for SDDPCM MPIO devices or disables the health check function of a device.

**Note:** Starting with SDDPCM v2123, when you issue this command to change the SDDPCM internal health check function for one or more devices from *enable* to *disable* or from *disable* to *enable*, the SDDPCM server daemon (pcmsrv) automatically starts or stops the health check function on the devices.

### Syntax:

```
▶▶ pcmpath set device num1 [num2] hc_interval t ▶▶
```

### Parameters:

*num1* [ *num2* ]

- When only *num1* is specified, the command applies to the hdisk specified by *num1*.
- When 2 device logical numbers are entered, this command applies to all the devices whose logical numbers fit within the range of the two device logical numbers.

*t* The range of supported values for health check interval is 1-3600 seconds. To disable the health check function of a device, set interval time to 0.

**Examples:** If you enter **pcmpath set device 2 10 hc\_interval 30**, the health check time interval of hdisk2 to hdisk10 is immediately changed to 30 seconds.

## pcmpath set device hc\_mode

The `pcmpath set device hc_mode` command dynamically changes the MPIIO device health check mode.

### Syntax:

```
▶▶ pcmpath set device num1 [num2] hc_mode option ▶▶
```

### Parameters:

*num1* [ *num2* ]

- When only *num1* is specified, the command applies to the hdisk specified by *num1*.
- When 2 device logical numbers are entered, this command applies to all the devices whose logical numbers fit within the range of the two device logical numbers.

*option*

Specifies one of the following policies:

- **enabled**, indicates the health check command will be sent to paths that are opened with a normal path mode.
- **failed**, indicates the health check command will be sent to paths that are in failed state.
- **nonactive**, indicates the health check command will be sent to paths that have no active I/O. This includes paths that are opened or in failed state.

**Examples:** If you enter `pcmpath set device 2 10 hc_mode enabled`, the health check mode of MPIIO hdisk2 to hdisk10 is immediately changed to the enabled mode.

## pcmpath set device path

The **pcmpath set device path** command sets the path of a device either to *online* or *offline*. You cannot remove the last path to an opened device from service. This prevents a data access failure from occurring.

**Attention:** If the device is closed, the **pcmpath set adapter offline** command does not have the last path offline protection.

### Syntax:

```
▶▶—pcmpath set device—device number—path—path-ID—online
offline————▶▶
```

### Parameters:

*device number*

The logical device number of the hdisk.

*path ID*

The path ID that you want to change, as displayed under “Path Name” by the **pcmpath query device** command.

*online*

Enables the path for service.

*offline*

Disables the path from service.

**Examples:** If you enter the **pcmpath set device 5 path 0 offline** command, path 0 for device 5 changes to Offline mode.

## Summary of command syntax

SDDPCM supports the following **pcmpath** commands:

- **pcmpath clear device** <n>/(<m> <n>) **count error/all**
- **pcmpath query adapter** [n]
- **pcmpath query adaptstats** [n]
- **pcmpath query device** [m n / -d <device\_model> -i
- **pcmpath query devstats** [m n / -d <device\_model>] -i
- **pcmpath set adapter** n **online | offline**
- **pcmpath set device** M **path** N **online | offline**
- **pcmpath set device** <n1> [n2] **algorithm** <option>
- **pcmpath set device** <n1> [n2] **hc\_interval** <t>
- **pcmpath set device** <n1> [n2] **hc\_mode** <option>
- **pcmpath disable port** <location> **ess** <essid>
- **pcmpath enable port** <location> **ess** <essid>
- **pcmpath open device** <m> **path** <n>
- **pcmpath query essmap**
- **pcmpath query portmap**
- **pcmpath query wwpn**
- **pcmpath query version**

**Note:** If the commands are used for a device, the *n* is the number of the device logical name. For example, **pcmpath query devstats 3** queries the device statistics for hdisk3.

If the commands are used for adapter, the *n* is the index of the adapter. For example, **pcmpath query adapter 2** queries the adapter statistics for the third adapter in adapter list order, which can be fscsi5.





---

# Index

## A

about this book vii  
accessing  
  AIX  
    add a data path volume to a  
      volume group SMIT panel 74  
    add a volume group with data  
      path devices SMIT panel 74  
    add paths to available data path  
      devices SMIT panel 73  
    backup a volume group with data  
      path devices SMIT panel 75  
    configure a defined data path  
      device SMIT panel 74  
    define and configure all data path  
      devices SMIT panel 73  
    display data path device adapter  
      status SMIT panel 73  
    display data path device  
      configuration SMIT panel 72  
    Display Data Path Device Status  
      SMIT panel 72  
    remake a volume group with data  
      path devices SMIT panel 75  
    remove a data path device SMIT  
      panel 74  
    Remove a Physical Volume from a  
      Volume Group SMIT panel 75  
adapters  
  firmware level 9, 97  
  LP7000E 4, 93  
  upgrading firmware level to  
    (sf320A9) 97  
adding  
  paths  
    AIX 36  
addpaths  
  command 68, 76  
  utility programs, AIX 76  
AIX  
  4.3.3  
    32-bit applications 10  
    64-bit applications 10  
  5.1.0  
    32-bit applications 10  
    64-bit applications 10  
  5.3.0  
    32-bit applications 10  
    64-bit applications 10  
  6.1.0  
    64-bit applications 10  
accessing  
  add a data path volume to a  
    volume group SMIT panel 74  
  add a volume group with data  
    path devices SMIT panel 74  
  add paths to available data path  
    devices SMIT panel 73  
  backup a volume group with data  
    path devices SMIT panel 75

AIX (*continued*)  
  accessing (*continued*)  
    configure a defined data path  
      device SMIT panel 74  
    define and configure all data path  
      devices SMIT panel 73  
    display data path device adapter  
      status SMIT panel 73  
    display data path device  
      configuration SMIT panel 72  
    Display Data Path Device Status  
      SMIT panel 72  
    remake a volume group with data  
      path devices SMIT panel 75  
    Remove a Physical Volume from a  
      Volume Group SMIT panel 75  
    removing a data path device SMIT  
      panel 74  
  applications  
    32-bit 10  
    64-bit 10  
  backing-up files belonging to an SDD  
    volume group 70  
  changing the path-selection policy 58  
  configuring  
    SDD 28, 33  
    volume group for failover  
      protection 63  
  dynamically adding paths to SDD  
    vpath devices of a volume  
      group 36  
  exporting  
    volume group with SDD 69  
  extending  
    an existing SDD volume group 70  
  fibre-channel  
    dynamic device tracking 59  
    requirements 6, 95  
  host attachment, installing 8  
  Hot Plug support 36  
  importing  
    volume group with SDD 68  
  installing  
    SDD 13  
    SDDPCM 98  
  migrating  
    an existing non-SDD volume  
      group to SDD vpath devices in  
      concurrent mode 85  
    non-SDD volume group to a SAN  
      Volume Controller SDD  
    multipath volume group in  
      concurrent mode 83  
    non-SDD volume group to an ESS  
      SDD multipath volume group in  
      concurrent mode 83  
  MPIO devices 111  
  NIM SPOT server 101  
  OS  
    upgrade 21

AIX (*continued*)  
  recovering  
    from mixed volume groups 69  
  removing  
    SDD from a host system 38  
    SDDPCM from a host system 109  
  restoring files belonging to an SDD  
    volume group 70  
  SDD utility programs 76  
  SDD-specific SMIT panels 71  
  the loss of a device path 66  
  trace 87, 118  
  unconfiguring SDD 34  
  unsupported environments 3, 92  
  upgrading 18  
  verifying SDD  
    configuration 35  
    installation 25  
AIX 5.1.0  
  32-bit 11  
  64-bit 11  
AIX 5.2.0  
  32-bit 11  
  64-bit 11  
AIX 5.3.0  
  32-bit 11  
  64-bit 11

## B

backing-up AIX files belonging to an  
  SDD volume group 70  
bootinfo -K command 11

## C

cfallvpath 41  
cfgmgr  
  command 8, 67, 68, 106  
  run *n* times where *n* represents the  
    number of paths per SDD  
    device. 68  
  run for each installed SCSI or  
    fibre-channel adapter 68  
changing  
  path-selection policy for AIX 58  
chdev command 66, 67  
command  
  addpaths 68, 76  
  bootinfo -K 11  
  cfgmgr 8, 67, 106  
  running *n* times for *n*-path  
    configurations 68  
chdev 66, 67  
datapath query device 63, 67  
datapath set device *N* policy  
  rr/fo/lb/df 59  
dpovgfix 63, 77  
dpovgfix vg-name 67  
extendvg 70

command (continued)

- extendvg4vp 70, 78
- hd2vp and vp2hd 76
- hd2vp vg\_name 19
- installp 7
- instfix -i | grep IY10201 7
- instfix -i | grep IY10994 7
- instfix -i | grep IY11245 7
- instfix -i | grep IY13736 7
- instfix -i | grep IY17902 7
- instfix -i | grep IY18070 7
- ls -al /unix 11
- lscfg -vl fcsN 9, 97
- lsdev -Cc disk 8, 107
- lsdev -Cc disk | grep 2105 28
- lsdev -Cc disk | grep SAN Volume Controller 28
- lslpp -l ibmSdd\_432.rte 25
- lslpp -l ibmSdd\_433.rte 25, 26
- lslpp -l ibmSdd\_510.rte 25, 26
- lslpp -l ibmSdd\_510nchacmp.rte 25, 26
- lspv 18, 65
- lsvg -p vg-name 65
- lsvgfs 18
- lsvpcfg 18, 62, 66, 77
- mkdev -l vpathN 34
- mksysb restore command 65
- mkvg 64
- mkvg4vp 64, 78
- odmget -q "name = ioaccess" CuAt 45
- pcmpath clear device count 124
- pcmpath disable ports 125
- pcmpath enable ports 127
- pcmpath open device path 129
- pcmpath query adapter 131
- pcmpath query adaptstats 132
- pcmpath query device 133
- pcmpath query devstats 136
- pcmpath query essmap 138
- pcmpath query portmap 139
- pcmpath query version 140
- pcmpath query wwpan 141
- pcmpath set adapter 142
- pcmpath set device 0 path 0 offline 146
- pcmpath set device algorithm 143
- pcmpath set device health\_check mode 145
- pcmpath set device path 146
- pcmpath set health\_check time interval 144
- querysn 40, 79
- restvg 70
- restvg4vp 70
- rmdev 67
- rmdev -dl dpo -R 18, 38, 109
- rmdev -dl fcsN -R 8, 107
- rmdev -l dpo -R 35
- savevg 70
- savevg4vp 70
- shutdown -rF 8, 107
- smitty 18
- smitty deinstall 7
- smitty device 18
- smitty uninstall 7

command (continued)

- table of, in installation package 11
- umount 18
- using 147
- varyoffvg 18, 28
- varyonvg vg\_name 19
- configuring
  - AIX
    - cabling storage side switch ports 8, 107
    - disk storage system 5
    - ESS 94
    - fibre-channel attached devices 6, 95
    - fibre-channel-attached devices 8, 106
    - SAN Volume Controller 6
    - volume group for failover protection 63
  - MPIO devices as the SAN boot device 116
  - SDD
    - AIX 33
    - AIX host 28
- conversion scripts
  - hd2vp 76
  - vp2hd 34, 76

## D

datapath
 

- query
  - device command 63, 67
  - set device N policy rr/fo/lb/df command 59

determining
 

- AIX
  - adapter firmware level 9, 97
- devices.fcp.disk.ibm.rte 2, 6
- devices.fcp.disk.ibm2105.rte 2
- devices.scsi.disk.ibm2105.rte 2

displaying
 

- AIX
  - ESS SDD vpath device configuration 62
- dpovgfix command 63, 77
- dpovgfix vg-name command 67
- dynamically opening an invalid or close\_dead path 59

dynamically removing or replacing adapters
 

- AIX Hot Plug support 36
- paths
  - AIX Hot Plug support 36
- dynamically removing paths 38
- dynamically replacing adapters
  - different type replacement 36
  - same type replacement 36

## E

ESS
 

- AIX
  - displaying SDD vpath device configuration 62
- devices (hdisks) 81

ESS (continued)
 

- LUNs 81
- exporting a volume group with SDD, AIX 69
- extending an existing SDD volume group, AIX 70
- extendvg command 70
- extendvg4vp command 70, 78

## F

failover protection
 

- AIX
  - creating a volume group from a single-path SDD vpath device 66
  - losing 65
  - manually deleted devices and running the configuration manager 67
  - side effect of running the disk change method 66
  - the loss of a device path 66
  - verifying load-balancing and failover protection 62
  - when it doesn't exist 63
- fibre-channel adapters
  - supported
    - AIX host systems 4, 93
- fibre-channel device drivers
  - configuring for AIX 6, 95
  - devices.common.IBM.fc 6, 96
  - devices.fcp.disk 6, 96
  - devices.pci.df1000f7 6, 96
  - installing for AIX 6, 95
  - supported
    - AIX host systems 4, 93

## H

HACMP
 

- concurrent mode 43
- enhanced concurrent mode 109
- hd2vp conversion script 46
- importing volume groups 46
- node fallover 55
- nonconcurrent mode 43
- persistent reserve 45
- recovering paths 55
- SDD persistent reserve attributes 45
- software support for enhanced concurrent mode 109
- software support for nonconcurrent mode 44
- hd2vp and vp2hd command 76
- hd2vp vg\_name command 19
- hdisk device
  - chdev 66
  - modify attributes 66
- healthcheck 113
- High Availability Cluster Multiprocessing (HACMP) 43
- host attachment
  - upgrade 21

## I

- IBM System p with static LPARs configured 10
- ibm2105.rte 5
- ibm2105.rte ESS package 4
- ibmSdd\_433.rte installation package
  - for SDD 1.2.2.0 removing 46
  - for SDD 1.3.2.0. SDD vpath devices unconfiguring 46
- importing a volume group with SDD, AIX 68
- installation package, AIX
  - devices.sdd.nm.rte 13
  - devices.sdd.43.rte 10, 34
  - devices.sdd.51.rte 10, 34
  - devices.sdd.52.rte 10
  - devices.sdd.53.rte 10
  - devices.sdd.61.rte 10
  - devices.sddpcm.52.rte 109
  - ibmSdd\_432.rte 25, 39, 43, 83
  - ibmSdd\_433.rte 25, 26, 39, 43, 45, 46, 83
  - ibmSdd\_510.rte 26, 39, 43
  - ibmSdd\_510nchacmp.rte 39, 43
  - SDD 13
- installing
  - AIX
    - fibre-channel device drivers 6, 95
    - host attachment 8
    - planning 1, 89
    - SDD 13
    - SDDPCM 98
- installp command 7
- instfix -i | grep IY10201 command 7
- instfix -i | grep IY10994 command 7
- instfix -i | grep IY11245 command 7
- instfix -i | grep IY13736 command 7
- instfix -i | grep IY17902 command 7
- instfix -i | grep IY18070 command 7
- invalid or close-dead path
  - dynamically opening 59

## K

- KB 137

## L

- load-balancing, AIX 62
- losing failover protection, AIX 65
- ls -al /unix command 11
- lscfg -vl fcsN command 9, 97
- lsdev -Cc disk | grep 2105 command 28
- lsdev -Cc disk | grep SAN Volume Controller command 28
- lsdev -Cc disk command 8, 107
- lslpp -l '\*Sdd\*' command 25
- lslpp -l ibmSdd\_432.rte command 25
- lslpp -l ibmSdd\_433.rte command 25, 26
- lslpp -l ibmSdd\_510.rte command 25, 26
- lslpp -l ibmSdd\_510nchacmp.rte command 25, 26
- lspv command 18, 65
- lsvg -p vg-name command 65
- lsvgfs command 18

- lsvpcfg 41
- lsvpcfg command 18, 62, 66, 77
- lsvpcfg utility programs, AIX 77

## M

- manual exclusion of disk storage system devices from the SDD configuration 40
- migrating
  - AIX
    - an existing non-SDD volume group to SDD vpath devices in concurrent mode 85
    - non-SDD volume group to a SAN Volume Controller SDD multipath volume group in concurrent mode 83
    - non-SDD volume group to an ESS SDD multipath volume group in concurrent mode 83
- mirroring logical volumes 84
- mkdev -l vpathN command 34
- mksysb restore command 65
- mkvg command 64
- mkvg4vp command 64, 78
- multipath SAN boot support 116

## N

- NIM SPOT server 101

## O

- odmget -q "name = ioaccess" CuAt command 45

## P

- path-selection policy
  - changing 59
  - failover only 58
  - load balancing 58
  - round robin 58
- pccpath
  - clear device count command 124
  - disable ports command 125
  - enable ports command 127
  - open device path command 129
  - query
    - adapter command 131
    - adaptstats command 132
    - device command 133
    - devstats command 136
    - essmap command 138
    - portmap command 139
    - set adapter command 142
    - version command 140
    - wwpn command 141
  - set device 0 path 0 offline command 146
  - set device algorithm 143
  - set device hc\_interval 144
  - set device health\_check mode 145
  - set device path command 146
- Persistent Reserve command set 45

- planning

- AIX

- adapter firmware level 9, 97
    - disk storage system 5
    - ESS 94
    - fibre-channel attached devices 6, 95
    - fibre-channel device drivers 6, 95
    - fibre-channel-attached devices 8, 106
    - installation 1, 89
    - preparing 5, 94
    - SAN Volume Controller 6
  - hardware requirements
    - fibre adapters and cables 92
    - supported storage devices 92
  - hardware requirements, AIX
    - disk storage systems 2
    - Fibre channel adapters and cables 2
    - Host system 2
    - SAN Volume Controller 2
    - SCSI adapters and cables 2
  - hardware requirements, SDDPCM
    - Fibre adapters and cables 92
    - Host system 92
    - supported storage devices 92
  - host system requirements, AIX 3
    - ESS 4
    - Fibre 4
    - SAN Volume Controller 4
    - SCSI 4
    - supported storage devices 93
  - host system requirements, SDDPCM
    - Fibre 93
  - software requirements, AIX
    - AIX operating system 2
    - ibm2105.rte ESS package 2
    - SCSI and fibre-channel device drivers 2
  - software requirements, AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1
    - fibre-channel device drivers 92
  - software requirements, SDDPCM
    - AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1 operating system 92
- preparing
  - AIX
    - SDD installation 5
    - SDDPCM installation 94
    - configure on AIX 28
- pvid 65, 83

## Q

- qdepth\_enable 32

## R

- recovering from mixed volume groups, AIX 69

- removing
  - SDD
    - from an AIX host 38
    - from an AIX host system 38
  - SDDPCM
    - from an AIX host 109
    - from an AIX host system 109
- replacing manually excluded devices in the SDD configuration 41
- requirements
  - hardware, AIX
    - disk storage systems 2
    - Fibre channel adapters and cables 2
    - Host system 2
    - SAN Volume Controller 2
    - SCSI adapters and cables 2
  - hardware, SDDPCM
    - fibre adapters and cables 92
    - Fibre adapters and cables 92
    - Host system 92
    - supported storage devices 92
  - host system, AIX 3
    - ESS 4
    - Fibre 4
    - SAN Volume Controller 4
    - SCSI 4
    - supported storage devices 93
  - host system, SDDPCM 93
    - Fibre 93
  - software
    - AIX operating system 2
    - ibm2105.rte ESS package, AIX 2
    - SCSI and fibre-channel device drivers, AIX 2
    - SDDPCM, AIX 5.2 TL07 (or later), AIX 5.3 TL03 (or later), or AIX 6.1 operating system 92
- restoring
  - AIX
    - files belonging to an SDD volume group 70
- restvg command 70
- restvg4vp command 70
- rmdev -dl dpo -R command 18, 38, 109
- rmdev -dl fcsN -R command 8, 107
- rmdev -l dpo -R command 35
- rmdev command 67

## S

- savevg command 70
- savevg4vp command 70
- SCSI adapter support
  - AIX host systems 4
- SCSI-3 Persistent Reserve command set 45
- SDD
  - adding paths
    - to vpath devices of a volume group 36
  - configuration
    - checking 35
  - host attachment, removing 40
  - installing
    - AIX 1

- SDD (*continued*)
  - server daemon
    - AIX host systems 12, 56
  - upgrading
    - automatically 14
    - manually 18
  - utility programs, AIX 76
  - verifying
    - configuration 35
    - vpath devices 81
- SDDPCM
  - installing
    - AIX 89
    - from AIX NIM SPOT server 101
  - maximum number of devices 104
  - path healthcheck time interval 113
  - pcmpath commands 122
  - server 107
  - server daemon 118
  - updating package 102, 104
  - verifying if the server has started 118
- sddserver.rte
  - AIX host systems 9
- sddsrv
  - AIX host systems 12, 56
  - for ESS Expert
    - AIX host systems 9
- secondary-system paging space, managing 61
- shutdown -rF command 8, 107
- SMIT
  - definition 13, 100
- smitty command 18
- smitty deinstall command 7
- smitty device command 18
- smitty uninstall command 7
- smitty, definition 13, 100
- synchronizing logical volumes 84
- System Management Interface Tool (SMIT) 100
  - definition 13, 100
  - using for configuring 28
  - using to access the Add a Data Path Volume to a Volume Group panel on AIX host 74
  - using to access the Add a Volume Group with Data Path Devices panel on AIX host 74
  - using to access the backup a Volume Group with Data Path Devices on AIX host 75
  - using to access the Configure a Defined Data Path Device panel on AIX host 73, 74
  - using to access the Define and Configure All Data Path Devices panel on AIX host 73
  - using to access the Display Data Path Device Adapter Status panel on AIX host 73
  - using to access the Display Data Path Device Configuration panel on AIX host 72
  - using to access the Display Data Path Device Status panel on AIX host 72

- System Management Interface Tool (SMIT) (*continued*)
  - using to access the Remake a Volume Group with Data Path Devices on AIX host 75
  - using to access the Remove a Data Path Device panel on AIX host 74
  - using to access the Remove a Physical Volume from a Volume Group panel on AIX host 75
  - using to back up a volume group with SDD on AIX host 70
  - using to backup a volume group with SDD on AIX host 75
  - using to create a volume group with SDD on AIX host 64
  - using to display the SAN Volume Controller SDD vpath device configuration on AIX host 62
  - using to display the SDD vpath device configuration on AIX host 62
  - using to export a volume group with SDD on AIX host 69
  - using to extend an existing SDD volume group on AIX host 70
  - using to import a volume group with SDD on AIX host 68
  - using to remove SDD from AIX host 38
  - using to remove SDDPCM from AIX host 109
  - using to restore a volume group with SDD on AIX host 75
  - using to restore a volume group with SDD vpath devices on AIX host 70
  - using to unconfigure SDD devices on AIX host 34
  - using to verify SDD configuration on AIX host 35

## U

- umount
  - command 18
- unconfiguring SDD on AIX 34
- unsupported environments
  - AIX 3, 92
- updating SDD using a PTF 19
- upgrade
  - AIX OS 21
  - host attachment 21
  - SDD packages 21
- upgrading
  - AIX
    - adapter firmware level 97
    - manually 18
  - SDD
    - automatically 14
  - SDD, manually
    - for AIX 4.3.2 18
    - for AIX 4.3.3 18
    - for AIX 5.1.0 18
    - for AIX 5.2.0 18
- using
  - AIX trace function 86

- using *(continued)*
  - ESS devices
    - directly, AIX 81
    - through AIX LVM 82
  - pcmpath commands 147
  - SAN Volume Controller devices
    - through AIX LVM 82
  - SDDPCM trace function, AIX 118
- utility programs, AIX
  - addpaths 76
  - dpovgfix 77
  - extendvg4vp 78
  - hd2vp and vp2hd 76
  - lsvpcfg 77
  - mkvg4vp 78
  - using disk storage system devices
    - through AIX LVM 82
  - using ESS devices directly 81
  - using SAN Volume Controller devices
    - through AIX LVM 82
  - using the SDDPCM trace function 118
  - using the trace function 86

## V

- varyoffvg command 18, 28
- varyonvg vg\_name command 19
- verifying
  - AIX
    - configuring SDD 35
    - SDD installation 25
- volume groups
  - AIX 64
  - mixed
    - dpovgfix vg-name 67
    - how to fix problem 67

## W

- Web sites
  - AIX APARs
    - maintenance level fixes and microcode updates 4
    - technology level fixes and microcode updates 93
  - information on the fibre-channel adapters that can be used on your AIX host 4, 93
  - information on the SCSI adapters that can attach to your AIX host 4









Printed in USA