



PART NUMBER
313482708

VERSION NUMBER
3.3

REVISION NUMBER
H

SVA PATH™

SOFTWARE

USER'S GUIDE
For AIX

PRODUCT TYPE
SOFTWARE

Shared Virtual Array SVA Path

Version 3.3

for AIX®

User's Guide

Part Number: 313482708

Information contained in this publication is subject to change without notice.

We welcome your feedback. Please contact the Global Learning Solutions Feedback System at:

GLSFS@Stortek.com

or

Global Learning Solutions
Storage Technology Corporation
One StorageTek Drive
Louisville, CO 80028-3256
USA

Please include the publication name, part number, and edition number in your correspondence if they are available.

Export Destination Control Statement

These commodities, technology or software were exported from the United States in accordance with the Export Administration Regulations. Diversion contrary to U.S. law is prohibited.

Restricted Rights

Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) and (2) of the Commercial Computer Software - Restricted Rights at FAR 52.227-19 (June 1987), as applicable.

Limitations on Warranties and Liability

Storage Technology Corporation cannot accept any responsibility for your use of the information in this document or for your use in any associated software program. You are responsible for backing up your data. You should be careful to ensure that your use of the information complies with all applicable laws, rules, and regulations of the jurisdictions in which it is used.

Proprietary Information Statement

The information in this document, including any associated software program, may not be reproduced, disclosed or distributed in any manner without the written consent of Storage Technology Corporation.

Should this publication be found, please return it to StorageTek, One StorageTek Drive, Louisville, CO 80028-5214, USA. Postage is guaranteed.

Edition F, May 2005, EC 132282

StorageTek and the StorageTek logo are trademarks or registered trademarks of Storage Technology Corporation. Other products and names mentioned herein are for identification purposes only and may be trademarks of their respective companies.

© Copyright Storage Technology Corporation 2001–2005. All rights reserved.

© Copyright Dot Hill Systems Corp. 2001. All rights reserved.

Table of Contents

Summary of Changes	vii
Preface	ix
Who Should Read This Guide	ix
Shared Virtual Array Documentation	ix
How to Obtain Software Documentation	ix
SVA Administrator for AIX Library	x
SVA Hardware Publications	x
Chapter 1. SVA Path Overview	1
Supported number of Paths and Devices	1
Build 741 and Lower	1
Build 742 and Higher	1
Failover/Failback Data Paths	2
Load Balancing	2
Dynamic Allocation of Device Resources	3
How SVA Path Works	3
Supported SVA Path Configurations	3
System Requirements	3
Summary of SVA Path Benefits	4
Document Overview	4
Chapter 2. SVA Path Hardware Setup	7
Fibre Addressing Concepts	7
Host Bus Adapters/Initiators	7
Domains in Fibre Channel Connection	7
Domain Numbering	8

Domain Numbering with Fibre and SVA Path	9
Hardware Preparation	9
Chapter 3. SVA Path Installation	11
Determining the SVA Path Build Number	11
On the CD-ROM Distribution	12
On Your System	12
Determining Your AIX Driver Version	12
Installing SVA Path on AIX	13
Installing SVA Path with SVA Administrator	15
Build 741 and Lower	15
Build 742 and Higher	16
Uninstalling SVA Path on AIX	16
Manually Saving SVA Path Configuration Files	17
SVA Path Device Naming in AIX	17
AIX Devices With Only One Path to the Host	18
Chapter 4. SVA Path Operation	21
The setsp Command	21
The spmon Command	24
Configuring Load Balancing	24
Tuning Path Balancing	27
How to Verify Load Balancing	30
Checking Path Changes	32
Adding FDEVs to an Existing Larger LUN	32
The sppath Command	32
Display Options for sppath	35
Ignoring and Reclaiming Devices with sppath	35
Understanding the setsp -a Screen	36
Changing the Configuration	38
Specifying a Device for setsp	39
Assigning a New Primary Path	39
Turning the Exclusion Setting Off and On	40
Redefining the Buffer Pointer Allocation	40
Turning Load Balancing On and Off	40

Changing the Retry Count and Retry Delay	41
Turning Failback Off and On	41
Dynamic Device Detection	41
Deleting Devices	42
Chapter 5. Diagnosing Errors.....	45
Troubleshooting AIX	45
Before Installing SVA Path 3.3	45
Failover Error Reporting in AIX	46
Recovering from an AIX Path Failure.....	46
Using PPRC with SVA Path	46
Error Messages	48
Appendix A. AIX Drivers and Firmware	51
HBA and Disk Drivers	51
HBA Firmware	52
System and Service Processor Firmware	52

Summary of Changes

EC	Date	Edition	Revision	Description
132282	May, 2005	F	H	See this edition for details.
132051	July, 2004	Fifth		See this edition for details.
128901	January, 2004	Fourth		See this edition for details.
128848	October, 2003	Third		See this edition for details.

Preface

This guide describes how to use the SVA Path Intelligent Data Path Management software. SVA Path provides improved performance and data accessibility for the StorageTek Shared Virtual Array (SVA).

Who Should Read This Guide

This guide is for data administrators, capacity planners, performance specialists, and system administrators. This guide assumes that you are familiar with Shared Virtual Array operations and AIX system administration.

Shared Virtual Array Documentation

This section lists software and hardware documentation for the Shared Virtual Array products.

How to Obtain Software Documentation

All of the Shared Virtual Array software publications are available from the following sources:

- On the SVA Software Publications CD-ROM (part number 3112953nn). To order a copy, contact StorageTek Publication Sales and Service at 800-436-5554 or send a fax to 303-661-7367.
- Online (for viewing and printing), at the StorageTek Customer Resource Center (CRC) Web site at:
www.support.storagetek.com.

Click on Software and go to the Shared Virtual Array Software list.

Note: Access to the CRC site requires a password. To obtain a password, call StorageTek Customer Support at 800-678-4430.

**SVA Administrator
for AIX Library**

- *Shared Virtual Array Administrator for AIX Command Quick Reference*
- *Shared Virtual Array Administrator for AIX Installation Guide*
- *Shared Virtual Array Administrator for AIX Messages*
- *Shared Virtual Array Administrator for AIX Quick Start Guide*
- *Shared Virtual Array Administrator for AIX User's Guide*

For any StorageTek software:

- *Requesting Help from Software Support*

**SVA Hardware
Publications**

Shared Virtual Array hardware publications are available from the following sources:

- On the SVA Hardware Publications CD-ROM (part number 3118447nn). To order a copy, contact StorageTek Publication Sales and Service at 800-436-5554 or send a fax to 303-661-7367.
- Online (for viewing and printing), at the StorageTek Customer Resource Center (CRC) Web site at:
www.support.storagetek.com.

Click on Disk Subsystems.

Note: Access to the CRC site requires a password. To obtain a password, call StorageTek Customer Support at 800-678-4430.

The V2X Shared Virtual Array (SVA) library consists of:

- *V2X Shared Virtual Array
Introduction*
- *V2X Shared Virtual Array
Operation and Recovery*
- *V2X Shared Virtual Array
Physical Planning*
- *V2X Shared Virtual Array
Planning, Implementation, and Usage*
- *V2X Shared Virtual Array
Reference*
- *V2X Shared Virtual Array
System Assurance*
- *Peer to Peer Remote Copy Configuration Guide*

The V960 library consists of:

- *V960 Shared Virtual Array
Introduction*
- *V960 Shared Virtual Array
Operation and Recovery*
- *V960 Shared Virtual Array
Physical Planning*
- *V960 Shared Virtual Array
Planning, Implementation, and Usage*
- *V960 Shared Virtual Array
Reference*
- *V960 Shared Virtual Array
System Assurance*
- *Peer to Peer Remote Copy Configuration Guide*

The 9500 SVA library consists of:

- *9500 Shared Virtual Array
Introduction*
- *9500 Shared Virtual Array
Operation and Recovery*
- *9500 Shared Virtual Array
Physical Planning*
- *9500 Shared Virtual Array
Planning, Implementation, and Usage*
- *9500 Shared Virtual Array
Reference*
- *9500 Shared Virtual Array
System Assurance*
- *Peer to Peer Remote Copy Configuration Guide*

Chapter 1. SVA Path Overview

This chapter provides an overview of SVA Path software and its features.

Supported number of Paths and Devices

SVA Path offers a new level of data accessibility and improved performance for the SVA. It eliminates the point of failure represented by a single input/output (I/O) path between servers and storage systems and permits I/O devices to be distributed across multiple paths.

Note: Starting from SVA Path 3.3 build 742, the number of supported paths has been enhanced. To determine the build level of your SVA Path distribution, see “Determining the SVA Path Build Number,” on page 11.

- Build 741 and Lower**
- 1024 FDEVs per LUN
 - 1024 LUNs (sum of all the LUNs handled by SVA Path)
 - 1024 paths (sum of all the paths handled by SVA Path)
 - 32 paths per LUN

This means that you can have 512 LUNs with two paths each, or up to 32 LUNs with 32 paths each.

- Build 742 and Higher**
- 1024 FDEVs per LUN
 - 1024 LUNs (sum of all the LUNs handled by SVA Path)

- 4096 paths (sum of all the paths handled by SVA Path)
- 32 paths per LUN

This means that you can have 128 LUNs with 32 paths each, or up to 1024 LUNs with 4 paths each.

Failover/Failback Data Paths

By providing alternate I/O paths from the server to the SVA, SVA Path provides uninterrupted access to mission-critical data. This substantially insulates server applications from I/O path failures.

In the event of a failed host bus adapter (HBA), interface cable, or channel I/O card within the SVA, SVA Path automatically reroutes I/O traffic to an alternate data path. Failover is essentially transparent, ensuring continuous access to data stored on the SVA. When configured in the recommended failback mode, SVA Path automatically restores the primary data path and system redundancy once the defective component is replaced.

Load Balancing

SVA Path supports up to 32 data paths between a host and any SVA logical device. While only two data paths are required for path failover capability, multiple data paths can be used to improve performance in one of the following ways:

1. by allowing SVA Path to uniformly distribute primary paths among all available I/O paths. This is the default behavior of SVA Path.
2. by manually assigning I/O traffic for a logical drive to a particular path. The administrator with an understanding of the I/O load patterns of his or her applications can optimize performance through an intelligent choice of paths.

| Dynamic Allocation of Device Resources

In SVA configurations with multiple servers attached to the same storage device, SVA Path allows the system administrator to assign a logical drive to one server and prevent the other servers in the SVA from accessing that same logical drive.

How SVA Path Works

SVA Path's filter driver resides between the file system drivers and the SCSI disk device driver. I/O requests are passed from the file system through SVA Path, then the SCSI disk driver and ultimately to the hardware.

SVA Path monitors the flow of I/O requests through the layered driver architecture. When it detects a failure along an I/O path, it automatically reroutes the request to an alternate path. Failover to the redundant I/O path is transparent to server applications and permits continuous access to the information stored on the disk array(s). To the operating system, there is only a slight delay in normal I/O operations during path failover; existing drive numbers and device access functions continue to work as expected.

If an I/O operation fails on all available paths, it is retried periodically on the primary path according to the retry delay and retry count parameters specified with `setsp`.

Supported SVA Path Configurations

SVA Path supports single-server configurations.

System Requirements

Before proceeding to the next chapter, please contact StorageTek Software Support to verify that your site meets the minimum hardware and software requirements supported by StorageTek. You will need to provide details about your site configuration, including:

- Host system processor and memory

- Host disk space
- Host operating system level
- SVA subsystem microcode
- Host bus attachments (HBAs)
- Switches
- High-availability software
- Third-party disk management software

Summary of SVA Path Benefits

- Increases potential subsystem throughput by directing I/O through multiple host adapters and SVA channels. Logical drives can be assigned to host bus adapters, manually balancing the I/O load across paths.
- Provides continuous access to mission-critical data by insulating server applications from I/O path failures.
- Installs easily and is transparent to server applications.
- Allows you to limit access to devices in a multiple server environment using Logical Unit Number (LUN) exclusion.

Document Overview

This manual describes how to install and configure SVA Path on systems running the AIX operating system.

- Chapter 2 explains configuring your hardware in preparation for installing SVA Path.
- Chapter 3 contains instructions for installing SVA Path.
- Chapter 4 explains configuring and operating SVA Path.
- Chapter 5 offers assistance in diagnosing error messages.

- Appendix A provides instructions on downloading and installing the latest AIX device drivers.

Note that user documentation for products used with SVA Path, including AIX documentation, is referenced throughout this manual. Have your hardware and operating system manuals available for reference.

Chapter 2. SVA Path Hardware Setup

Fibre Addressing Concepts

Host Bus Adapters/Initiators

The terms “host bus adapter” and “initiator” mean essentially the same thing. Typically, the HBA is a card within the host that, in its role as initiator, issues commands on the Fibre channel.

Domains in Fibre Channel Connection

StorageTek uses the concept of “domains” to allow open systems hosts access to blocks of logical devices (the domains) within an SVA. A domain is an additional layer of device addressing, but one that is manually configured by the StorageTek Customer Service Engineer (CSE) in the SVA. This layer of addressing divides the SVA into “domains of access.” There can be up to 16 (0–15) domains per SVA, with each domain having one target with 256 LUNs. There is a limit of 1024 (4096 for V2X) total devices available within an SVA.¹

Each open systems host initiator is connected with Fibre cables to a controller card port, giving it access to the devices that have been configured within its domain. (An open systems host cannot see devices in domains other than the one to which it is attached.)

Full SVA Path functionality requires that redundant initiators can access the SVA over redundant data paths.

1. Using all allowed domains, targets, and LUNs, there are more than 1024 logical devices, but the SVA has a limit of 1024 logical devices. This limit is not applicable to V2X

Figure 2-1 shows two data paths connecting the open systems platform to the attached SVA using Fibre cables.




Figure 2-1 Fibre Paths from the Host to the SVA

Note: In the above figure, a domain can represent up to 256 logical devices (1 Target \times 256 LUNs = 256 Logical Devices). There is a limit of 1024 logical devices with an SVA (4096 for V2X).

Domain Numbering

Using Fibre connection, the domain number is configured at the SVA operator panel by the StorageTek Customer Service Engineer.

The domain number is never seen by the open systems host; from its point of view, just a target and logical unit number are involved in an I/O operation.

Domain Numbering with Fibre and SVA Path

Without SVA Path, it is not advisable to have more than one path from a single host set to the same domain number on a given SVA. Two hosts accessing the same domain could attempt to share the same LUN and corrupt the data stored on it.

SVA Path addresses the problems of such a setup through its LUN exclusion feature, which allows association of each LUN of the domain to only one host.

With SVA Path installed, the system needs to be set up with two or more paths accessing the same domain for the failover/failback feature to be enabled.

Note: With only one path, failover/failback cannot be enabled, but the system could still benefit from the LUN exclusion feature in a multi-host environment.

A single host without SVA Path installed, accessing the same domain through different paths, would see different devices representing the same LUN and could be confused by it.

Hardware Preparation

When adding SVA Path for AIX to an existing, configured SVA subsystem, no configuration changes are required.

Chapter 3. SVA Path Installation

Before beginning the installation, please contact StorageTek Software Support to verify that your site meets the minimum hardware and software requirements supported by StorageTek. You will need to provide details about your site configuration, including:

- Host system processor and memory
- Host disk space
- Host operating system level
- SVA subsystem microcode
- Host bus attachments (HBAs)
- Switches
- High-availability software
- Third-party disk management software

Note: If there is an earlier version of SVA Path installed on your system, older than that of your distribution CD-ROM package, remove it before continuing with this installation. See “Installing SVA Path on AIX,” on page 13.

Determining the SVA Path Build Number

To determine if an older version of SVA Path is installed on your system, compare the build numbers as described below.

On the CD-ROM Distribution

To determine the build number of your SVA Path distribution, perform the following steps:

1. Insert the CD-ROM.
2. Log on as root.
3. Type the following commands:

```
#/usr/sbin/mount -v cdrfs -o ro  
/dev/cd0/mountpoint
```

```
#ls mountpoint
```

For builds 752 and after you will see the following subdirectories:

- aix4
- aix5
- hacmp
- svadisk

For builds prior to 752 you will see the SVA Path package name, SVAPath.aix-3.3.build*nnnn*, where *nnnn* is the build number.

On Your System

To determine the build number of SVA Path installed on your system, examine the output of the command:

```
#lslpp -L |grep SVAPath
```

You will see the following output:

Fileset	Level	State	Type	Description (Uninstaller)
SVAPath	3.3.290.0	C	F	SVAPath Software (build 0 <i>nnn</i>)

where *nnn* is the build number.

Determining Your AIX Driver Version

Before proceeding with SVA Path installation, you should verify that the latest versions of the AIX device drivers are installed on

your system. See Appendix A., “AIX Drivers and Firmware” for complete details.

Installing SVA Path on AIX

Follow these steps to install the SVA Path driver and its supporting files.

Note: Before proceeding, refer to the ReadMe file in the */mountpoint* directory for information on any software or procedure changes that may have occurred after this manual was printed.

1. Set up your system hardware according to the instructions in their respective user manuals and in “SVA Path Hardware Setup,” on page 7.
2. Before installing SVA Path software, the host must be able to see and access all storage devices through all available paths. For example, if you have two HBAs and redundant paths to your storage system, each logical drive should be displayed twice when viewing your list of available devices.

If you cannot see all available storage devices through every path, verify that you have the latest fibre-channel or SCSI HBA drivers loaded on your system.

3. To install the software, insert the CD-ROM.
4. Log on as `root`.
5. Type the following command:

```
#/usr/sbin/mount -v cdrfs -o ro /dev/cd0  
/mountpoint
```

Note: Installation of SVA Path will rebuild your system kernel. You will need to reboot in order to use SVA Path devices.

6. Issue an `ls` command to locate the install filename and build level for use in subsequent steps.

7. Depending on your build level, type the appropriate command.

– For builds 752 and after, type:

```
#cd /mountpoint/aix4 or
```

```
#cd /mountpoint/aix5
```

– For builds prior to 752, type:

```
#cd /mountpoint
```

8. Type the following command:

```
#!/svapath-install install SVAPath.aixn-3.3  
.buildnnnn
```

where *n* is the SVA Path for AIX version number (such as 4 or 5), and *nnnn* is the build number from Step 6. For example, SVAPath.aix5-3.3.build0752.

9. Follow the on-screen instructions to completion.

Note: Starting with SVA Path 3.2, build 720, if the `spd.conf`, `sppath.conf`, and `spmon.conf` files are backed up to `/opt/storagetek/svap_bak`, they are restored automatically during this SVA Path install.

The current installed `*.conf` files are copied to `/opt/storagetek/svap_install`.

The `stk.renumber fileset` is executed.

If you need a fresh configuration of SVA Path, you may do the following:

```
#cp /opt/storagetek/svap_install/spd.conf  
/etc/spd.conf
```

```
#cp /opt/storagetek/svap_install/spmon.conf  
/etc/spmon.conf
```

```
#cp /opt/storagetek/svapth_instll/  
sppath.conf /etc/sppath.conf
```

10. Reboot the system:

```
shutdown -Fr
```

11. Run the following command to check whether all available devices are seen by SVA Path.

```
#setsp -a
```

12. For the first SVA Path installation, run the following command to configure the devices:

```
#setsp -e0 -l all
```

13. If the BULL AIX SCSI-2 device driver is used on the host, issue the following commands:

```
#setsp -T -l all  
#chdev -a pv=yes -l hdiskhdiskN
```

Run the following command to start all spd devices:

```
#setsp -S -l all
```

Installing SVA Path with SVA Administrator

SVAA uses a designated LUN on the SVA for administrative commands. This LUN is the Extended Control and Monitoring (ECAM) device.

Since patch 13, SVAA has the capability of handling several ECAM devices per subsystem; this provides a basic failover mechanism. Consult your SVAA documentation for details on running SVAA's `sibconfig` command to add several ECAM devices

Note: Starting with Build 742, SVA Path is fully compatible with SVAA. According to the build number of SVA Path, read the instructions in the following chapters, in order to determine the best way of managing SVAA ECAM devices. To determine the build level of your distribution, see “Determining the SVA Path Build Number,” on page 11.

Build 741 and Lower To avoid any interaction between SVAA and SVA Path, it is advised to remove all ECAM devices from the control of

SVA Path. After SVA Path is installed and the host has rebooted, perform the following steps:

1. Type the following:

```
setsp -a
```
2. Record the `spd` number and paths displayed on the screen above the ECAM device path.
3. For all `spd` numbers found, run the following commands from a Command Prompt window:
 - a. Instruct SVA Path to terminate the `spd` driver:

```
#setsp -T -lspdN
```
 - b. Instruct SVA Path to ignore these paths/disks:

```
#sppath -I cXtYdZ
```
 - c. Once all the ECAM devices are removed from SVA Path control, make the update permanent:

```
#setsp -g
```

Changes will take effect only after the next reboot of the host.

Build 742 and Higher

SVA Path has been improved to handle ECAM devices. To fully benefit from SVA Path functionality, it is recommended to put all ECAM devices under SVA Path control. Once SVA Path is installed, no unique actions need to be performed for ECAM devices; proceed as for any other devices, to include them under SVA Path control.

Uninstalling SVA Path on AIX

Note: For SVA Path 3.2, Build 720 or later, the uninstall sequence will automatically save existing SVA Path configuration files. These configuration files will be used when the SVA Path upgrade is installed. For versions of SVA Path prior to Build 720, configuration files must be saved manually when uninstalling SVA Path (See “Manually Saving SVA Path Configuration Files,”

on page 17). Manually saved files will then be automatically restored during installation of upgrade code (Build 720 or later).

1. Mount the CD-ROM as described in the installation section, and run the following command from the CD directory:

```
#./svapath_install uninstall
```
2. Restore any application-specific files that were modified during the installation procedure.
3. Restore the hardware configuration.
4. Reboot the system.

Manually Saving SVA Path Configuration Files

If you are uninstalling SVA Path in order to upgrade from a version of SVA Path before 3.2, Build 720, you may want to save and restore your current configuration files.

Perform the following commands to manually back up existing configuration files:

```
#setsp -g  
  
#mkdir /opt/storagetek/svap_bak  
  
#cp /etc/spd.conf  
/opt/storagetek/svap_bak/spd.conf  
  
#cp /etc/sppath.conf  
/opt/storagetek/svap_bak/sppath.conf  
  
#cp /etc/spmon.conf  
/opt/storagetek/svap_bak/spmon.conf
```

These saved backup files will be restored automatically when the upgrade is installed.

SVA Path Device Naming in AIX

Disk device filenames in AIX use the following format:

```
/dev/hdiskX  
/dev/rhdiskX
```

where X = the controller instance number. A sample listing follows:

```
hdisk2 Available 10-70-01 Other FC SCSI Disk Drive
hdisk3 Available 10-70-01 Other FC SCSI Disk Drive
hdisk4 Available 10-70-01 Other FC SCSI Disk Drive
hdisk5 Available 20-60-01 Other FC SCSI Disk Drive
hdisk6 Available 20-60-01 Other FC SCSI Disk Drive
hdisk7 Available 20-60-01 Other FC SCSI Disk Drive
```

In a typical, single-path disk storage system, each physical disk is represented in the host's /dev directory by two special device files representing the character and block mode device instances. For simplicity, we reference the set of device files as a single device filename and disregard the first part of the pathnames that distinguish between raw and block modes.

As a layered device driver that runs on top of the standard AIX disk driver (hdisk), SVA Path creates its own device files. The SVAPath created special files are only used for SVA Path internal processing. From an application view, the SVA Path created files are not used. The AIX logical volume manager and applications continue to use the hdisk and associated identifiers.

When we have multiple path, for each LUN, we have as many devices (and thus special files) on the host. If we consider that all these file that give access to the same LUN form a group, applications continue to use one special file of such groups: the smallest numbered one. This can be checked with the output of 'setsp -a'. Access to the other special files in the group is blocked by SVAPath.

Note: It is the virtual device file that will be used by applications to access that device. The original, redundant data paths should never be used to access the device, or the data it contains could be corrupted.

**AIX Devices With
Only One Path to the
Host**

SVA Path creates virtual device files for all supported devices and blocks access to them through their original device filenames whether or not they are accessible via redundant paths. Virtual

device files are created for devices that have only one physical path for either of the following reasons:

- A device could actually have more than one path, but only one was functional at the time that SVA Path was installed.
- The device could be part of a SAN configuration where it is necessary to prohibit applications on the local host from accessing the device because the device was assigned to another host on the SAN.

Note: Boot devices do not qualify as supported devices, and virtual device files are never created for them. (For a complete description of qualifying devices, refer to “The sspath Command,” on page 32.)

Table 3-1 *Installed SVA Path Files (Partial List^a)*

File	Description
/usr/lib/drivers/spd	SVA Path driver
/etc/sspath.conf	sspath configuration file (installed empty)
/etc/spmon.conf	spmon configuration file
/etc/spd.conf	spd configuration file (installed empty)
/etc/spd/bin/sspath	qualifies and claims SVA storage devices for SVA Path control
/etc/spd/bin/setsp	configures SVA Path parameters
/etc/spd/bin/spmon	monitors device paths and implements load balancing

a. A full list of all files installed is provided during the installation sequence, or by running the `ls1pp -f SVAPath` command.

Chapter 4. SVA Path Operation

This chapter describes SVA Path commands and configuration options.

There are three basic commands in SVA Path

- `setsp` is used to examine and configure the system's operating parameters, and as such is the command most often invoked by the user.
- `spmon` monitors path states and implements load balancing.
- `sppath` identifies devices to be placed under SVA Path's control.

The `setsp` Command

The `setsp` command is used for most configuration tasks.

Table 4-1, “`setsp` Command Options”, on the next page gives a brief explanation of `setsp` command options. Those used to configure SVA Path device parameters are covered in some detail in “Changing the Configuration”, beginning on page 38. In these cases, the pages containing more detailed explanations of command options are noted parenthetically.

The output of `setsp -a` is also used in conjunction with operating system events to determine the nature and physical location of failures. This is covered in “Diagnosing Errors”, starting on page 45.

Command options that take arguments (shown in braces after the command) *require* an argument and should not be run without one.

Table 4-1 *setsp Command Options*

Option	Effect
-a	show current device configuration and state
-b{0 1}	set load balancing for a logical unit. (pages 37 and 40). Must be used with -l -b0 disables load balancing for the device; -b1 enables it Note: SVA Path performs load balancing at the whole-LUN level only. In the case of larger LUNs, it does not perform load balancing on the individual FDIDs within a larger LUN.
-d<n>	set a retry delay of <i>n</i> for a logical drive, where <i>n</i> is the interval between retries in milliseconds (page 38 and 41). Must be used with -l
-e{0 1}	set exclusion for a logical drive (pages 37 and 40). Must be used with -l -e0 includes the device; -e1 excludes it; excluded devices are not accessible by user applications; devices are excluded by default
-f{0 1}	set failback for a logical drive. Must be used with -l -f0 disables failback for the device; -f1 enables it (pages 38 and 41)
-g	generate configuration files after an <code>sppath</code> command (Page 34)
-i	show contents of driver configuration files

Table 4-1 setsp Command Options

Option	Effect
-L <i>parameter</i>	show current device configuration according to condition(s) defined by setsp command option parameter(s) (e.g., setsp -L -e1 lists all excluded devices; setsp -L -b0 -f1 lists all devices that have load balancing disabled and failback enabled); acceptable parameters are: -l, -e, -p, -b, -r, -d, and -f
-l{<n> all}	specify a logical drive for the command, where; <i>n</i> is the drive's spd number; all specifies all devices (page 39)
-N	runs a command to change the configuration files without affecting the running system (changes will take effect at the next boot)
-n<n>	allocate <i>n</i> buffer pointers for a logical drive. Must be used with -l. Note: This value is always set to "1"; any other value will be ignored.
-p<n>	select primary path <i>n</i> . Must be used with -l.
-r<n>	set a retry count of <i>n</i> for a logical drive, where <i>n</i> is the number of times a command will be retried (pages 38 and 41)
-S	start the spd driver
-T	terminate the spd driver
-u{0 1 2}	show devices by their configuration status: -u0 shows all available devices; -u1 shows configured disks; -u2 shows unconfigured disks

Table 4-1 *setsp Command Options*

Option	Effect
-v	runs a command in verbose mode
-x	Remove an spd device from the configuration file. The spd driver will no longer configure it after a reboot. (<code>setsp -g -l</code> reconfigures the spd, and returns it to SVA Path control.) The result is similar to ignoring a device (using <code>sppath -I</code>).

The spmon Command

The `spmon` command is primarily associated with load monitoring and balancing of SVA Path functional devices (FDEVs).

The terms “device” and “FDEV” refer to a logical disk (LUN) as viewed by the host operating system, the applications, and the users. An FDEV can emulate one of a variety of SCSI and count-key-data (CKD) disk devices. A LUN may be comprised of a single FDID, or it may be a larger LUN comprised of many FDIDs.

Note: SVA Path performs load balancing at the whole-LUN level only; it does not perform load balancing on individual FDIDs within a larger LUN.

Note: For optimum performance when striping data across multiple LUNs, disable load balancing on the devices being striped using the `setsp -b0` option. Also, the primary path of each LUN should be assigned to a different HBA in a way that will ensure static balance of the I/O.

Configuring Load Balancing

Note: If you have modified the primary paths manually and you wish to keep your setup configuration, you should not activate the load balancing option. If you do, `spmon` may automatically

reassign the paths in which case you will need to reboot or manually reassign all paths to retrieve your setup configuration.

Load balancing is enabled or disabled using `setsp -l all -b0` (disable) or `setsp -l all -b1` (enable) commands, with the default setting being load balancing is enabled. Parameters for load balancing are stored in the configuration file `spmon.conf`. This file is read automatically after each modification at the end of the expiration of the last measurement interval.

Following is a listing of available parameters in `spmon.conf`:

- `balance-threshold percentage`

This parameter represents the maximum difference in load between the highest-loaded and lowest-loaded paths, as a percentage of the highest load, before the paths are considered balanced. At 100%, paths will never be considered unbalanced; at 50%, imbalance occurs when the load on the least-busy path is less than half of that of the busiest.

A value of 10% specifies that any difference of more than 10% between the most and least busy paths triggers the path balancing algorithm.

- `reassignment-threshold percentage`

The `reassignment-threshold` parameter limits reassignments that do not reduce imbalance enough to be worthwhile. When considering a path reassignment, the load balancing algorithm computes a target value for the load it wants to transfer from the busiest path to the least busy path. The reassignment threshold is the percentage by which a less-than-optimal move may deviate from the target value and still be considered a candidate. For example, a value of 99% means that reassignments may be considered even if they result in only a very small reduction in the imbalance, while a value of 50% specifies that a device shall be reassigned only if such a reassignment will reduce the imbalance to less than 50% of what it was before the move. Lower percentage values discourage path reassignments.

- `measurement-interval time`

This parameter accepts a positive integer value with a suffix of “s” (seconds), “m” (minutes), or “h” (hours). A reasonable minimum value will be based on the CPU load presented to the system by the algorithm and the maximum value is based on the amount of time that can pass before the driver's internal counters overflow.

- `reassignment-limit number`

This parameter specifies the maximum number of devices that should be moved in one pass of the algorithm. The default value is equal to one half of the FDEVs.

- `read-overhead ms-per-cmd ms-per-sector`

This parameter is used to specify how bus connect time overhead is estimated for read commands. The first value specifies the number of microseconds estimated for read command overhead, while the second is an estimate of connect time required for each 512 bytes of data requested. For each read operation, the sum of the command overhead and the product of the transfer length and the per-sector overhead is added to a counter that is used to estimate overall bus utilization on a per-FDEV and per-channel basis.

- `write-overhead ms-per-cmd ms-per-sector`

This parameter is analogous to the read-overhead statement. Sample values for these two statements are:

```
read-overhead 1000 120
write-overhead 1000 160
```

- `log-data-directory directory-name`

This parameter specifies that log files should be placed in the specified directory each time the load balancing algorithm runs. The data files in the directory are named using the “path group name,” consisting of the names of the HBAs in each group concatenated with hyphens, followed by `-fdevs.csv` or `-paths.csv`. These files contain the data

used as input to the load balancing algorithm. The `-fdevs` file contains the following fields:

```
SecondSinceProgramStarted,  
FdevName,  
CurLoadMilliSec,  
oldPath,  
NewPath
```

The `-paths.csv` file contains these fields:

```
SecondSinceProgramStarted,  
PathNum,  
PathName,  
CurLoadMilliSec
```

If the directory name contains a space or ends in a backslash, enclose the entire directory name in double quotes.

Tuning Path Balancing

The path balancing algorithm captures details about the I/O load going to each device, then enters an algorithm which performs several tests to determine whether a device movement will improve the balance across all possible physical paths.

For the purpose of this section, moving a device means re-directing I/O from one physical interface to another. This creates a certain amount of work to be done by the SVA. While tuning using the parameters below, you should consider how to achieve a satisfactory balance while moving the least amount of devices:

Parameters should be changed one at a time and then monitored for a period of time. An indication of an incorrect parameter value can either be no device movements or too many device movements. Try to always err on the side of no device movements and adjust slowly until device movement is seen.

Additionally, the `spmon` command can be a valuable asset to observe system load numbers either real-time, or over a period of

time. The show option of spmon will display load information every time it is issued.

There is also a log-data-directory parameter in the spmon.conf file which identifies a file to place load information every time the algorithm runs, i.e., every measurement interval. This file can then be imported into a graphing utility, and load peaks and lows can be seen over time.

Read / Write μ s-per-cmd & sector

The first two parameters which can be used to tune the algorithm are:

- Read-overhead μ s-per-cmd μ s-per-sector
- Write-overhead μ s-per-cmd μ s-per-sector

These parameters aid in determining the I/O load to the devices. Since writes require more subsystem resources than reads, these parameters set the skew. They must be altered to match the prevalent RFA blocksize being used by the operating system per the chart below:

Table 4-2Blocksize

	512 Byte	2K Byte	4K Byte	8K Byte	16K Byte
Read μ s-per-sector	110	45	25	25	24
Write μ s-per-sector	180	68	44	33	32

Note: When tuning, leave the Read / Write μ s-per-command at 1000. It is also recommended to only change one of these parameters since they have a co-relationship.

Measurement Interval

The next parameter to consider tuning is the measurement interval time. This parameter determines how long the I/O load will be measured before entering the algorithm itself.

Raising the value too high could overrun the program buffers and cause errors.

Setting the value too low renders the sample unique, to the extreme where it is not representative of the actual load.

The value of this parameter should remain in the range of 20 - 60 seconds.

Balance Threshold Percentage

Balance Threshold Percentage is the first actual Load Balancing Algorithm parameter. It determines how far apart, in load, the paths can be allowed to get before the algorithm progresses with further tests. The actual formula is:

$$\frac{\text{High_Load_Path} - \text{Low_Load_Path}}{\text{High_Load_Path}} > \text{Balance_Threshold_Percent}$$

If the result of this formula is true, the algorithm will continue to the Re-assignment Threshold test. (See below.)

The default value for Balance Threshold Percentage is 10%. Try setting it lower if the load is mostly small I/Os and higher if there are mostly large I/Os. As with all parameters, an indication of an incorrect parameter is either no movement or too many movements. Bear in mind that with a smaller number, more movement is likely; larger numbers restrict movement.

Re-assignment Threshold Percentage

Re-assignment Threshold Percentage is the second Load Balancing Algorithm parameter, and determines whether a device is worth moving. It forces the software to locate the best device to move and therefore achieve the tolerable balance set in the Balance Threshold Percent parameter. It asks the question, "Will an X% improvement be made to the balance?"

The default value is 50%. Setting this parameter too high can cause too many movements and impact device performance. Too low, and no devices will ever qualify to be moved.

For example, if there were a theoretical load on path A of 100 and a load on path B of 50:

If Balance Threshold Percent = 10%, the maximum tolerable out-of-balance condition is within 65 - 85, with the ideal at 75.

If the Re-Assignment Threshold Percentage is set to 50%, only a device with a load of 12.5 to 25 would be a candidate to be moved.

12.5 would change path A to 87.5 and path B to 62.5: a 50% improvement.

Re-Assignment Limit Number

The final tuning parameter is Re-Assignment Limit Number. This parameter controls how many devices can be moved after each Measurement Interval.

Set this parameter to 1/2 the number of paths under SVA Path control on that host, for example:

- If there are 2 paths, set it to 1
- If there are 4 paths, set it to 2
- If there are 3 paths, experiment with both 1 and 2, starting with 1.

How to Verify Load Balancing

In Figure 4-1, “Example setsp -a Output in Determining Load Balancing”, there are three SVA Path devices; spd0, spd1 and spd2.

```
setsp -a
```

spd	Path/disk	Status	Pri	Exc	Buf	Balance	RtrCnt	RtrDly	FailBack
0	c6t1d0/9	Good	X		1	1	30	10000	1
	c7t1d0/11	Good							
spd0 = hdisk9						ID = "STK V960 0024000200000000"			
1	c6t1d1/3	Good	X		1	1	30	10000	1
	c7t1d1/4	Good							
spd1 = hdisk3						ID = "STK V960 0024004800000000"			
2	c6t1d2/10	Good	X		1	1	30	10000	1
	c7t1d2/12	Good							
spd2 = hdisk10						ID = "STK V960 0024004900000000"			

Figure 4-1 Example setsp -a Output in Determining Load Balancing

- Status shows the current state of the path. Good paths are functioning normally. Bad paths have failed. Excluded paths are unavailable to applications on this host.

In our example (Figure 4-2), the command `spmon show` will show that each device in Figure 4-1 is accessible via two paths, using HBAs “fscsi1” and “fscsi2”:

```

Total I/O load across monitored HBAs = 0ms

  HBA   Device      I/O Load  %load
-----
fscsi1 fscsi1         3ms    100.0%
fscsi2 fscsi2         0ms     00.0%

Device spd0 [STK V960 00240002000000000]
FDev   Path              I/O Load  HBA Load
-----
spd0:0  0/hdisk11           1ms        3ms    33.3%

Device spd1 [STK V960 00240048000000000]
FDev   Path              I/O Load  HBA Load
-----
spd1:0  0/hdisk3            1ms        3ms    33.3%

Device spd2 [STK V960 00240049000000000]
FDev   Path              I/O Load  HBA Load
-----
spd2:0  0/hdisk10           1ms        3ms    33.3%

```

Figure 4-2 Example spmon show Output

The column headings describe the various fields on the screen:

- **FDEV** (Functional Device) identifies the SVA Path device and the zero-based index of the FDEV within the device, separated by a colon.
- **Path** identifies the current path assigned to the FDEV and the name of the host bus adapter used by that path.
- **I/O Load** estimates channel utilization time during the current measurement interval. The I/O Load may be expressed as one of the following formats:

milliseconds (with an “ms” suffix) if it is less than one second.

seconds (with an “s” suffix) if it is less than 300 seconds.

hours:minutes:seconds (hh:mm:ss) if it is more than 59 minutes.

- **HBA Load** identifies the I/O load contributed by traffic to/from this FDEV as a percentage of all I/O though the current HBA used to access this FDEV. The figure in the right-hand column is a ratio I/O Load/HBA Load, expressed as a percentage.

In this example, `fscsil` is carrying 100% of the load. At the next measurement interval, `spmon` will automatically calculate and redistribute the I/O load between the two paths, significantly reducing the maximum load on `fscsil`.

Checking Path Changes

To log path change messages issued from the `spmon` daemon, type the following line at the end of the `/etc/syslog.conf` file:

```
daemon.info    /<logfile name>
```

For more information on how to log other system events, refer to the `/etc/syslog.conf` file and the `syslog` man pages.

Adding FDEVs to an Existing Larger LUN

Since load balancing is performed at the whole-LUN level only, all FDEVs belonging to a larger LUN are assigned to the same path, regardless of when they are added to the LUN.

The `sppath` Command

The `sppath` command is run automatically at boot after SVA Path is installed. `sppath`'s main function is to create the configuration file that identifies devices to be put under SVA Path's control. Its command options are used to display or modify this device set.

Note: Altering the physical connections to the SVA requires that the host be rebooted for SVA Path to recognize the new configuration.

sppath examines disk devices attached to the system to determine whether any physical devices are accessible via redundant paths and whether those devices should be put under the control of the SVA Path driver. Qualifying devices are written to the `sppath.conf` file. This file should not normally be modified directly by the user.

To qualify, devices must *not* be boot devices and must have the appropriate inquiry data.

The vendor ID must match one of:

- STK
- IBM
- RSBA

The product ID must match one of:

- 9200
- 9393
- 9500
- V960
- V2X

sppath's command options are used to display or modify this device set.

Table 4-3, "sppath Command Options" gives a synopsis of sppath options. They are explained in greater detail below.

Table 4-3 *sppath Command Options*

Option	Effect
-d	display debug information
-D	clear the list of ignored devices. Should be followed by <code>setsp -g</code> (see note below)
-I {cX[tY[dZ]]}	ignore the device or group of devices specified; devices are specified in the form cX, cXtY, or cXtYdZ. All devices that match will be ignored; Should be followed by <code>setsp -g</code> (see note below) Example: <code>sppath -I c5</code> will ignore all devices under controller 5; <code>sppath -I c5t0</code> will ignore all devices under target 0 of controller 5.
-v	display the contents of <code>/etc/sppath.conf</code> after writing the file

Note: After running `sppath` with options `-I` or `-D`, you must run `setsp -g` in order for the changes to be reflected in SVA Path's configuration file `spd.conf`. While most changes are immediate, *all* changes take effect at the next system reboot.

Display Options for sspath

The `-v` option causes `sspath` to display the contents of `sspath.conf` after updating it (Figure 4-3, “Using `sspath` in Verbose Mode”).

```
SPD=0 c6t1d0 dev=42, 7 type=2 SANID="STK V960 0024000200000000"  
SPD=0 c6t1d0 dev=42, 6 type=2 SANID="STK V960 0024000200000000"  
SPD=1 c6t1d1 dev=42, 1 type=2 SANID="STK V960 0024004800000000"  
SPD=1 c6t1d1 dev=42, 2 type=2 SANID="STK V960 0024004800000000"  
SPD=2 c6t1d2 dev=42, 8 type=2 SANID="STK V960 0024004900000000"  
SPD=2 c6t1d2 dev=42, 9 type=2 SANID="STK V960 0024004900000000"
```

Figure 4-3 Using `sspath` in Verbose Mode

The `-d` option displays inquiry data in raw format and is generally used only for debugging.

Ignoring and Reclaiming Devices with `sspath`

The `-I` option accepts symbolic device names corresponding to controllers or specific disks and omits them from `sspath.conf`.

This prevents them from being put under SVA Path’s control.

If a device is to be ignored, all of its paths should be specified with `-I` options. Multiple devices can be specified in a single `sspath` command, but each device specified must be preceded by `-I`. Once specified, the ignored device is remembered in `sspath.conf` and will be ignored until the list of ignored devices is cleared with `setsp` using the `-D` option. Ignored devices will be displayed as UNKNOWN by `setsp -a`. The ignored devices will behave like standard disk drives not under the control of SVA Path.

Run the following commands to ignore a device:

```
#setsp -T -l<spdN>  
#sspath -I <cXtYdZ>  
#setsp -g
```

Alternately, you can also run `sspath -I`, followed by a reboot of the host, in order for the operating system to recognize the ignored device.

Note: The `sspath -I` (ignore) command should not be confused with the `setsp -e` (exclude) command. The

former removes the device completely from SVA Path's control, treating it exactly as though it is incompatible with SVA Path and could not be recognized and claimed. The `setsp -e` command is intended chiefly for multiple host configurations and prevents particular logical devices from being accessed by a host.

The `-D` option clears the entire list of ignored devices, allowing any eligible device to be placed under SVA Path's control upon the next reconfiguration, either from a manually-entered `cfgmgr` command or a system reboot.

Understanding the `setsp -a` Screen

Use the `setsp -a` command to display the default path configurations.

```
setsp -a
```

spd	Path/disk	Status	Pri	Exc	Buf	Balance	RtrCnt	RtrDly	FailBack
0	c6t1d0/9 c7t1d0/11	Good Good	X		1	1	30	10000	1
spd0 = hdisk9							ID = "STK V960 0024000200000000"		
1	c6t1d1/3 c7t1d1/4	Good Good	X		1	1	30	10000	1
spd1 = hdisk3							ID = "STK V960 0024004800000000"		
2	c6t1d2/10 c7t1d2/12	Good Good	X		1	1	30	10000	1
spd2 = hdisk10							ID = "STK V960 0024004900000000"		

Figure 4-4 Sample Output of `setsp -a` Command

The column headings identify the various fields on this screen, the last seven of which are user-configurable parameters for the device. The `setsp -a` output fields (with their default values, when applicable) are described below. The commands used to change the default settings are described in the following section.

- `spd` is the SVA Path driver number, an ID assigned to the device by SVA Path and the name of the `spd` special device file created by SVA Path to access the storage. The disk

number is given as `hdiskx`, where *x* is a number used by AIX ODM.

- `Path/disk` shows the device names and disk numbers for each of the redundant physical paths to the device. Their appearance in the `setsp -a` output facilitates interpreting these events in terms of the `spd` device names by which applications access devices. This field shows the disk number in the form `cXtYdZ`.
- `Status` shows the current state of the path. Good paths are functioning normally. Bad paths have failed. Excluded paths are unavailable to applications on this host.
- `Primary` shows which of the physical I/O paths connecting the device to the host's host bus adapters (or controllers) is defined as primary (marked by an X). Initially, primary path assignments are distributed evenly among the available paths. Figure 4-4, "Sample Output of `setsp -a` Command" shows this as an alternating pattern in a dual-path configuration: `spd0` has a primary path to controller 7, `spd1` uses its path to controller 8 as primary, and `spd2` alternates back to controller 7. The device's duplicate path(s) are not used unless the primary path fails or the load balancing option is selected for that device.
- `Exclude` indicates the exclusion setting, which is used to keep particular servers from seeing particular logical drives. As a safety measure, SVA Path excludes all devices from host access by default, giving them an exclusion setting of 1. Excluded devices are marked by an X in this column.
- `Buf` is the number of buffer pointers (or buffer structures) pre-allocated for each logical device. Values between 1 and 100 are permitted; however, AIX can have only one buffer pointer. The default value is 1.
- `Balance` indicates whether dynamic load balancing is enabled for the device. When load balancing is enabled, SVA Path tracks the volume of I/O on each path and periodically reassigns the LUN's path as needed to keep the I/O load balanced across all HBAs. 1 means load balancing

enabled. 0 means load balancing not enabled. Load balancing is enabled by default for SVA devices.

- `RtrCnt` (retry count) is the number of times a failed I/O will be retried on the primary path after it has tried unsuccessfully to use its alternate path(s) and returned again to the primary. When the specified number of retries have failed, the I/O fails. The default value is 30. The highest value allowed is 100.

Note: When all paths for the `spd` are “Bad,” the `RtrCnt` value is reset to 0 until a “good” path is detected. Then the `RtrCnt` value will be reset to the configured value.

- `RtrDly` (retry delay) is the time interval, in milliseconds, between the retry attempts described in the preceding parameter. The default value is 10000 ms. This value can not be set above 100,000 ms (100s).
- `FailBack` indicates whether failback is enabled for the logical device. When failback is enabled (the default setting of 1), SVA Path will keep testing a path that has failed and return it to service (as the primary path or in sequential load balancing) as soon as the path has been restored.

If the default configuration is satisfactory, no reconfiguration of SVA Path is necessary.

In addition to the information defined by the `setsp -a` output's column headings, the last part of the last line for each device (`ID=`) shows the physical device's unique identifier, which is derived from the device's inquiry data.

Changing the Configuration

The information displayed by `setsp -a` is stored in the configuration file `spd.conf`, which should *never* be edited directly. User-configurable parameters must be changed exclusively through the `setsp` command options provided for that purpose and are described in more detail in the following sections.

Specifying a Device for setsp

To name a specific device in any `setsp` command, use the `-l` option (note that this is a lowercase letter L, and not the number 1) followed by the device's `spd` number, as expressed in the syntax illustrations used in the remainder of this chapter as `-l<x>`. To name all devices under SVA Path's control, use `-l all`. A device's `spd` number is listed in the first column of the `setsp -a` output (See `setsp -a` example on page 30).

Whether you are applying the configuration command to a single device or to all the devices, only one parameter can be changed per command.

Assigning a New Primary Path

The syntax for changing an SVA LUN's primary data path is

```
setsp -l<x> -p<n>
```

where `<x>` is the `spd` number (or `all`) and `<n>` is the number of the new path. The path number can be obtained by simply counting down `setsp`'s list of physical paths (in the `Path/disk` column), starting from zero.

The command to change the primary path shown in Figure 4-4, on page 36, for the SVA device with the `spd` ID of 0 from its default path through controller 7 to the path through controller 8, then, would be:

```
#setsp -l0 -p1
```

If you run a `setsp` command with the verbose option (`-v`), the configuration change will be displayed on screen (below).

```
=====
spd   Path/disk           Status   Pri Exc Buf Balance RtrCnt  RtrDly FailBack
=====
    0   c6t1d0/0X          Good     X    1   1    30   10000    1
spd0 = hdisk0                               ID = "STK V960 0024000200000000"
=====
```

Figure 4-5 Using `setsp -p` in Verbose Mode

If an SVA LUN comprises multiple FDEVs, all FDEVs are assigned to the specified path.

Turning the Exclusion Setting Off and On

The syntax for changing the exclusion setting for a device is

```
setsp -l<x> -e{0|1}
```

where <x> is the spd number (or all) and the -e option takes one of two arguments:

- -e0 turns exclusion off (makes the device visible to the host);
- -e1 turns exclusion on (excludes the device).

In a single-host configuration all devices should be visible to the host. Devices can be included (or unexcluded) by using the command.

```
#setsp -l all -e0
```

All devices may be included, because no other host can access them at the same time. However, in a multiple-host environment, where all spds are visible to SVA Path on all hosts, spds must be either excluded or included so that a host shares no spds. Use the ID number under the spd column on the `setsp -a` output to identify devices. Do not use the `spdX` number.

Note: When including a device you must run the following command:

```
#cfgmgr
```

Redefining the Buffer Pointer Allocation

The syntax for changing the number of buffer pointers pre-allocated for a given device is:

```
setsp -l<x> -n<n>
```

where <x> is the spd number (see “Specifying a Device for setsp” on page 39) and <n> is the new value.

Note: The value of <n> is always set to “1”; any other value will be ignored.

Turning Load Balancing On and Off

The syntax for changing a device's load balancing mode is

```
setsp -l<x> -b{0|1}
```

where $\langle x \rangle$ is the spd number (or all) and the `-b` option takes one of two arguments:

- `-b0` turns load balancing off;
- `-b1` turns load balancing on.

Changing the Retry Count and Retry Delay

The syntax for changing the retry count for a device is

```
setsp -l $\langle x \rangle$  -r $\langle n \rangle$ 
```

where $\langle x \rangle$ is the numerical element of the spd number (see “Specifying a Device for setsp” on page 39) and $\langle n \rangle$ is the number of times a failed I/O will be retried on the primary path (after its alternate paths have been tried unsuccessfully) before the path is marked as failed (with a Status of Bad).

The syntax for changing the retry delay for a device is

```
setsp -l $\langle x \rangle$  -d $\langle n \rangle$ 
```

where $\langle x \rangle$ is the numerical element of the spd number and $\langle n \rangle$ is the interval between the retries specified by the retry count parameter.

Turning Failback Off and On

The syntax for changing the a device’s failback mode is

```
setsp -l $\langle x \rangle$  -f{0|1}
```

where $\langle x \rangle$ is the spd number and the `-f` option takes one of two arguments:

- `-f0` turns failback off;
- `-f1` turns failback on.

Dynamic Device Detection

You may be able to add new storage devices and place them under SVA Path's control without requiring a reboot of the host. This is also known as dynamic LUN allocation.

To add a fibre channel device to an existing path, perform the following steps after the devices are physically connected:

1. Reconfigure the host:

```
#cfgmgr
```

Note: When you run `cfgmgr`, the following commands are automatically run to instruct SVA Path to detect the new devices (`sppath`), create an updated configuration file (`setsp -g`), start SVA Path on the new device (`setsp -S -l<new_spdN>`).

2. To obtain the new `spd` number, run:

```
#setsp -a
```

3. Unexclude the device:

```
#setsp -e0 -l<new_spdN>
```

4. If the BULL AIX SCSI-2 device driver is used on the host, issue the following commands:

```
#setsp -T -l all
```

```
#chdev -a pv=yes -l hdisk<hdiskN>
```

Repeat this for all newly discovered `hdisks`

```
#setsp -S -l all
```

Deleting Devices

Note: Prior to deleting a device from SVA Path, you must remove the device from the associated volume group and logical volume configuration, as necessary.

Use the following steps to delete a device:

1. From the command prompt, run the following command to instruct SVA Path to terminate the `spd` driver:

```
#setsp -T -l(n)
```

2. From the command prompt, run the following command to instruct SVA Path to ignore these paths/disks:

```
#sppath -I cxydz
```

3. From the command prompt, run the following command to make the update permanent:

```
#setsp -g
```

4. Use the SVAA CLI or the SVAC to delete the device.

Note: To remove a specific device from the ignored devices list, first run the `sppath -D` command to remove *all* devices from the list, then run the `sppath -I cxydz` command for each device you wish to return to the list.

Chapter 5. Diagnosing Errors

Troubleshooting AIX

Before Installing SVA Path 3.3

Before installing SVA Path software, the host must be able to see and access all storage devices through all available paths. For example, if you have two HBAs and redundant paths to your storage system, each logical drive should be displayed twice when viewing your list of available devices.

- **If you can't see all available storage devices** through every path, verify that you have the latest fibre channel or SCSI HBA drivers loaded on your system.
- **To load HBA drivers in AIX**, run the `installp` command or use the SMIT utility.

Once the proper driver is loaded, the system will be able to communicate with your HBA. All devices on the bus will be listed under the HBA within the Devices menu. If the host is unable to see the HBA and/or its devices, please upgrade to the latest driver. Most drivers can be obtained from the HBA manufacturer's web site.

If the latest driver is installed, please review the hardware configuration attached to the host. See “Hardware Preparation,” on page 9 for examples.

Failover Error Reporting in AIX

The output from `setsp -a` (Figure 5-1) now shows that every device's path 0 (those corresponding to device filenames beginning c5) have a Status designation of Bad.

```
=====
spd  Path/disk          Status  Pri Exc Buf Balance RtrCnt  RtrDly FailBack
=====
  0  c5t1d0/17          Bad     X   1   1   30   10000   1
     c6t1d0/12          Good
spd0 = hdisk12          ID = "STK V960 0024000200000000"
=====
  1  c5t1d1/13          Bad     X   1   1   30   10000   1
     c6t1d1/22          Good
spd1 = hdisk13          ID = "STK V960 0024004800000000"
=====
  2  c5t1d2/19          Bad     X   1   1   30   10000   1
     c6t1d2/14          Good
spd2 = hdisk14          ID = "STK V960 0024004900000000"
=====
```

Figure 5-1 `setsp -a` Output after the Failed Path Is Detected

Recovering from an AIX Path Failure

Path failures generate a lot of output, resulting primarily from the disk driver's attempts at error recovery.

The output of `setsp -a` will return to normal (Figure 5-2).

```
=====
spd  Path/disk          Status  Pri Exc Buf Balance RtrCnt  RtrDly FailBack
=====
  0  c5t1d0/17          Good    X   1   1   30   10000   1
     c6t1d0/12          Good
spd0 = hdisk12          ID = "STK V960 0024000200000000"
=====
  1  c5t1d1/13          Good    X   1   1   30   10000   1
     c6t1d1/22          Good
spd1 = hdisk13          ID = "STK V960 0024004800000000"
=====
  2  c5t1d2/19          Good    X   1   1   30   10000   1
     c6t1d2/14          Good
spd2 = hdisk14          ID = "STK V960 0024004900000000"
=====
```

Figure 5-2 `setsp -a` Output after the Failed Path Is Back Online

Using PPRC with SVA Path

As the LUNs on the secondary PPRC subsystem are R/W locked, SVA Path will generate error messages and consider all the paths to these LUNs as bad. To avoid this, proceed as follows prior to creating PPRC pairs (`-addpairs`):

1. Unmount the file systems and vary off the related LUN volume groups that will be used to create PPRC pairs.
2. Run this command for SVA Path to terminate the spd driver:

```
#setsp -T -l<spdN>
```

3. From the command prompt, issue the following command for SVA Path to ignore the spd drivers corresponding to the PPRC pairs:

```
#sppath -I <cXtYdZ>
```

4. Run this command to make this update permanent:

```
#setsp -g
```

If you plan to recover or delete pairs (-deletpairs) *and* you wish to see the LUNs through SVA Path, proceed as follows:

1. Clear the list of ignored paths:

```
#sppath -D
```

2. Re-generate the configuration:

```
#sppath -g
```

3. Start the spd drivers:

```
#setsp -S
```

Error Messages

Following are error messages you may see in the spd error log (/etc/spd/bin/spd.msg).

Note: Only the “Error Number” and “Message” fields are included in the log.

Error Number	Message	Severity	Explanation
6800	spd device started.	Information	spd devices have been started by a system reboot or a <code>setsp -S</code> command.
6801	not enough memory.	Warning	Cannot allocate memory in the kernel for a given set of devices.
6802	disk excluded by this host.	Information	The disk cannot be accessed by this host.
6803	use <code>setsp</code> to change exclusion status.	Information	Not used.
6804	found a conflict path.	Warning	There is a mismatch between configured and running spds. A possible cause is running <code>setsp -g</code> which attributes new spds to existing hdisks. Starting the new spds while the old ones associated with the hdisks are still running would generate this error.
6805	device inquiry data mismatch.	Warning	There is a mismatch in device configuration. A possible cause is deleting a LUN from the SVA subsystem and re-creating it with a different FDID, requiring a change of spd. Attempting to start an spd without rebooting or reconfiguring the host would generate this error.

Error Number	Message	Severity	Explanation
6806	no good paths found.	Warning	spd has failed to start.
6807	path is opened already.	Warning	Upon starting an spd, an exclusive path could not be opened. A possible cause is that a volume group has been varied on already.
6808	have a bad path(s).	Warning	spd has failed to start.
6809	can't detach.	Warning	Not used. spd has failed to stop.
680A	path failover.	Critical/ Information	A path failover has occurred. Possible causes: <ol style="list-style-type: none"> 1. An I/O failure has occurred on one path to a device while there is at least one other good path to the device. You need to investigate and fix the bad path. 2. If a device has load-balancing turned off, when the primary path status is changed from "bad" to "good," I/O will be redirected to the primary path.
680B	path ok.	Information	Path status has changed from "bad" to "good."
680C	path error.	Critical	Error while accessing a path.
680D	spd device stopped.	Information	An spd device has been stopped. If an spd is not stopped, no message is written to the log or to the output of the <code>setsp -T</code> command.

Error Number	Message	Severity	Explanation
680E	vdev_size changed.	Information	The size of a LUN/spd device has changed. The size stored in the kernel memory is different from the size derived from inquiry data.
680F	spd info.	Information	Not used.
6810	all paths are bad.	Critical	All paths for a device are bad. In an HACMP environment, when this occurs on the HACMP node and there is at least one other node still active, the AIX <code>halt -q</code> command is executed to shut down the bad node.

Appendix A. AIX Drivers and Firmware

This appendix provides instructions for downloading and installing the latest AIX device drivers and related firmware.

HBA and Disk Drivers

To display the drivers currently installed on your system, type the following:

```
lslpp -L devices.pci.df1* devices.fcp.disk.rte  
devices.scsi.disk.rte
```

To download the latest drivers, go to the following URL:

```
https://techsupport.services.ibm.com/server/aix.fdc
```

Fill in the search form as follows:

- In the OS level box, select the AIX level running on your system.
- In the Search by box, select Fileset or PTF number.
- In the Search String field, type either of the following:
 - `devices.pci.df1*` to display HBA drivers
 - `devices.fcp.disk.rte` or `devices.scsi.disk.rte` to display disk drivers

After downloading the drivers, you can install them using `smitty install_latest` from the download directory.

HBA Firmware

It is also necessary to verify that the HBAs (6227, 6228) are running the most current levels of firmware (microcode). This is independent of the AIX version on the host.

To display the HBA firmware level, type the following:

```
lscfg -vl fcsX
```

where *X* is the number for your adapter (0,1,2,3....).

If you are not at the minimum required levels, you may experience problems. You must download and install the latest IBM-approved versions from the IBM site, which is at the following URL:

```
http://techsupport.services.ibm.com/server/mdownload/
```

This site includes instructions for unbundling and installing the firmware. You should print and follow these instructions.

Note: The password to unbundle the drivers is: ****RS/6000****.

System and Service Processor Firmware

If your system does not boot after you upgrade the HBA microcode versions, your system and/or service processor firmware are probably down level. These need to be updated as well.

You can download the system/service processor firmware and associated upgrade instructions from the following URL:

```
http://techsupport.services.ibm.com/server/mdownload/
```

You will need to remove the HBAs from the machine, upgrade the firmware, and then re-install the HBAs.

