



PART NUMBER
313433606

VERSION NUMBER
3.3

REVISION NUMBER
F

SVA PATH™

SOFTWARE

USER'S GUIDE
For Windows

PRODUCT TYPE
SOFTWARE

Shared Virtual Array SVA Path

Version 3.3

for Windows

User's Guide

Part Number: 313433606

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 E, 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
Related SVA Software Publications	x
SVA Hardware Publications	xi
Chapter 1. SVA Path Overview	1
Failover/Failback Data Paths	1
Load Balancing	2
Dynamic Allocation of Device Resources	2
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	5
Fibre Addressing Concepts	5
Host Bus Adapters/Initiators	5
Domains in Fibre Channel Connection	5
Domain Numbering	6
Domain Numbering with Fibre and SVA Path	7
Hardware Preparation	7

Chapter 3. SVA Path Installation	9
Installing SVA Path on Windows 2000 or Windows 2003	9
Installing SVA Path with SVA Administrator	12
Uninstalling SVA Path	13
Manually Saving SVA Path Configuration Files	14
Installed Files	14
SVA Path Device Names on Windows 2000/2003	15
Chapter 4. SVA Path Operation	19
The setsp Command	19
The spmon Command	22
Configuring Load Balancing	22
Tuning Path Balancing	25
How to Verify Load Balancing	28
Adding FDevs to an Existing Lun	30
The sppath Command	30
Display Options for sppath	32
Ignoring and Reclaiming Devices with sppath	33
Understanding the setsp -a Screen	34
Changing the Configuration	36
Specifying a Device for setsp	36
Assigning a New Primary Path	37
Turning the Exclusion Setting Off and On	37
“Unknown Devices” Errors in Windows 2003	38
Redefining the Buffer Pointer Allocation	39
Removing LUNs from SVA Path	39
Excluding a Device Using the -x Option	39
Turning Load Balancing On and Off	40
Changing the Retry Count and Retry Delay	40
Turning Failback Off and On	40
Dynamic Device Detection	41
Deleting Devices	42
Chapter 5. Diagnosing Errors	43
Errors in the Event Viewer	43
VDS Errors in the System Event Log	45

Appendix A. Configuring SVA Path With Microsoft Cluster Service	47
Minimum System Requirements	47
Environment Description	48
Configuring Microsoft Cluster Service	50
Hardware Configuration	50
Preparing the Devices	50
Microsoft Cluster Service Installation	52
Microsoft Cluster Service Default Resources	52
Installing and Configuring SVA Path	53
Using SVA Path in a Microsoft Cluster Service Environment	56
Adding New Volumes	56
Deleting Existing Volumes	58

Summary of Changes

EC	Date	Edition	Revision	Description
132282	May, 2005	E	F	See this edition for details.
128848	October, 2003	Fourth		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 Windows 2000, and/or Windows 2003 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) website at: www.support.storageitek.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.

**Related SVA
Software
Publications**

**For Shared Virtual Array Administrator for open systems
platforms (AIX, HP-UX, Linux, Solaris, and Windows):**

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

For SVA Administrator for S/390:

- *Shared Virtual Array Administrator for OS/390
Configuration and Administration*
- *Shared Virtual Array Administrator for OS/390
Installation, Customization, and Maintenance*
- *Shared Virtual Array Administrator for OS/390
Messages and Codes*
- *Shared Virtual Array Administrator for OS/390
Reporting*

For SVA SnapShot for S/390:

- *SVA SnapShot for OS/390
Installation, Customization, and Maintenance*
- *SVA SnapShot for OS/390
User's Guide*

For SVA Console for Windows (SVAC):

- *Shared Virtual Array Console for Windows
Quick Start 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) website 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 SVA 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*

Chapter 1. SVA Path Overview

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

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.

SVA Path supports up to:

- 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.

SVA Path supports Microsoft Cluster Service in a Windows 2000 or Windows 2003 environment. See Appendix A., “Configuring SVA Path With Microsoft Cluster Service” on page 47.

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 three 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.
3. by enabling automatic load balancing. In this mode of operation SVA Path monitors the load on each path and reassigns LUNs/FDevs to balance the load across all available 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 multi-initiator (SVA) environment using LUN (logical unit number) exclusion.

Document Overview

This manual describes how to install and configure SVA Path on systems running the Windows 2000 or Windows 2003 operating systems.

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

Note that user documentation for products used with SVA Path, including Microsoft Windows 2000 and Windows 2003 documentation, is referenced throughout this manual. Have your hardware and operating system manuals available for quick 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 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 total devices available within an SVA.(4096 for V2X)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. (4096 for 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

It is not normally advisable to have more than one path from a single host set to the same domain number to a given SVA (in case two hosts attempt to share the same LUN and so corrupt the data stored on it). SVA Path requires exactly this configuration for failover to work.

SVA Path manages multiple paths from a single host, using identical domain numbers.

Hardware Preparation

When adding SVA Path for Windows to an existing, configured SVA/SVAA system, no configuration changes are required.

Chapter 3. SVA Path Installation

This chapter describes how to install SVA Path for use with an SVA.

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

Installing SVA Path on Windows 2000 or Windows 2003

Note: If a previous version of SVA Path is installed on your system, remove it before continuing with this installation (See “Uninstalling SVA Path,” on page 13).

SVA Path is distributed as the `setup.exe` file in the root directory of the CD-ROM. Follow these steps to install the SVA Path driver and its supporting files.

1. Domain addresses are automatically configured to be set to zero (refer to Chapter 2., “SVA Path Hardware Setup”, beginning on page 5).
2. Log in as `Administrator` on the host on which you are installing SVA Path.
3. Before installing SVA Path, it is imperative that the host is 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 by the **Disk Management** program. If you cannot see all available storage devices through every path, verify that you have the latest fibre channel HBA drivers loaded on your system.

In Windows 2000/2003, new devices are detected at boot time, prompting a wizard that will guide you through choosing and installing an HBA driver. If a new device is not automatically detected by Windows 2000/2003, or if the wizard fails to install a driver for any reason, click on the **Add/Remove Hardware** icon under the **Control Panel** and follow the onscreen directions. All devices on the bus will be listed under the HBA within the **Device Manager** menu, which can be accessed by clicking on **Control Panel**, then **Administrative Tools**, then **Computer Management**. 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.

4. Put the CD-ROM containing the SVA Path distribution files in the drive from which you want to install.
5. Click the `Start` button and Select `Run`. In the `Run` dialog box, choose the CD-ROM drive, and navigate to the `root` directory.

6. Double-click on the `setup.exe` program and follow the step-by-step instructions during the installation process.
7. Before proceeding, refer to the `README` file in the root directory for information on any changes to the software or installation procedure that may have occurred after this manual was printed.
8. If this is a first-time installation of SVA Path, continue with Step 9.

If this is an upgrade, and you are restoring saved SVA Path configuration files, make the following entries:

```
>copy C:\tmp\sva_bak\spd.conf
<SVAPath_install_dir>\spd.conf

>copy C:\tmp\sva_bak\sppath.conf
<SVAPath_install_dir>\sppath.conf

>copy C:\tmp\sva_bak\spmon.conf
<SVAPath_install_dir>\spmon.conf
```

9. Reboot the host.
10. Open a Command Prompt and verify that all available drives are seen by SVA Path by running the following command from the installation directory:

```
>sppath -v
```

11. Run the following command to save and implement changes to configuration files:

```
>setsp -g
```

Note: To chose individual devices to be claimed by SVA Path, use `-l` in the next step, as described in “Specifying a Device for setsp” on page 36.

12. SVA Path excludes all devices from host access by default. To turn the exclusion setting off for all devices and make all SVA Path-compatible devices accessible by the host, run the command:

```
>setsp -e0 -l all
```

Device exclusion is covered in more detail on pages 35 and 37.

13. Reboot the host.
14. Verify that all available drives are seen by SVA Path by running the command:

```
>setsp -a
```

15. Edit any other application-specific files to reflect the new device names. New device files— identical to the pre-installation device files, except for their controller numbers—are generated during SVA Path installation for all SVA devices (including those accessible by only a single path). Any applications already configured to use the older device files to access SVA LUNs must be reconfigured to use the new pathnames (an example of this is given in the instructions below for installing SVA Path with SVA Administrator).

The installation is now complete.

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.

Note: Since patch 13, SVAA has the possibility to handle several ECAM devices per subsystem. This provides a fail over mechanism but, to avoid any interaction between SVAA and SVAP, it is advised to remove all ECAM devices from the control of SVA Path. Consult your SVAA documentation for details on running SVAA's sibconfig command to add several ECAM devices

After SVA Path is installed and the host has rebooted, run `setsp -a`, record the spd number and paths displayed on the screen, above the ECAM device driver letter. Then, for all the spd

number found, run the following commands from a Command Prompt window:

1. Instruct SVA Path to terminate the spd driver:

```
>setsp -T -l<spdN>
```
2. Instruct SVA Path to ignore these paths/disks:

```
>sppath -I <cXbYtZdN>
```
3. Once all the ECAM devices removed from SVA Path control, make the update permanent::

```
>setsp -g
```

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

Uninstalling SVA Path

Note: Before uninstalling SVA Path, existing SVA Path configuration files must be manually saved if they are to be re-installed with an SVA Path upgrade. (See “Manually Saving SVA Path Configuration Files,” on page 14).

Saved configuration files must be manually restored during installation of SVA Path upgrade code.

To uninstall SVA Path:

1. Under Control Panel, open:
Add/Remove Software.
2. Select SVA Path 3.3 and chose:
Change/Remove.
Select "Yes" when asked to confirm the removal of the software.
3. Reboot the host.

4. Restore any application-specific files that were modified during the installation procedure.
5. Restore the hardware configuration.
6. 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 0720, you may want to save and restore your current configuration files.

Perform the following commands to manually backup existing configuration files:

```
>setsp -g  
  
>mkdir C:\tmp\sva_bak  
  
>copy <SVAPath_install_dir>\spd.conf  
C:\tmp\sva_bak\spd.conf  
  
>copy <SVAPath_install_dir>\sppath.conf  
C:\tmp\sva_bak\sppath.conf  
  
>copy <SVAPath_install_dir>\spmon.conf  
C:\tmp\sva_bak\spmon.conf
```

These saved files must be manually restored when the upgrade (SVA Path version 3.3, build 0720) is installed.

Installed Files

During installation, the files listed in Table 3-1, “Installed SVA Path Files” are placed in your system. These files will be placed in whichever directory is selected at the time of install (with the exception of spd.sys). The default install directory will be:

```
\Program Files\Storage Tek\SVA Path 3.3
```

All of these files will be removed if SVA Path is uninstalled.

Table 3-1 Installed SVA Path Files

File	Description
\WINNT\System32\drivers\spd.sys	SVA Path driver
sppath.conf	sppath configuration file (installed empty)
spmon.conf	spmon configuration file
spd.conf	spd configuration file (installed empty)
sppath.exe	qualifies and claims SVA storage devices for SVA Path control
setsp.exe	configures SVA Path parameters
spmon.exe	monitors device paths and implements load balancing
Uninst.isu	uninstall script used by Install Shield

SVA Path Device Names on Windows 2000/2003

To display the current configuration of SVA Path devices, run the command:

```
>setsp -a
```

As shown in Figure 3-2, disk device filenames are listed as *cWbXtYdZ*. This naming convention correlates to the device structure in the Windows Registry, where *cW* represents Scsi Port *W*, *bX* represents SCSI Bus *X*, *tY* represents Target Id *Y* and *dZ* represents LUN *Z*.

```

setsp -a
=====
      spd      Path/disk      Status  Pri  ExcBufBalanceRtrCntRtrDlyFailBack
=====
      0         c3b0t1d0/6    GoodX3202030001
              c4b0t0d0/12    Good
HardDisk 6  I:          ID = "STK 9500 0000000010390000"
=====
      1         c3b0t1d1/7    ExcludedX3202030001
              c4b0t0d1/13    ExcludedXX
HardDisk 7          ID = "STK 9500 0000000010390001"
=====
      2         c3b0t1d2/8    GoodX3202030001
              c4b0t0d2/14    GoodX
HardDisk 8          ID = "STK 9500 0000000010390002"
=====

```

Figure 3-2 Output of setsp -a Showing Device Names

By interpreting the Windows naming convention into `setsp -a` output, SVA Path does not create new device names. Disk Management sees only the logical drives as described in the Registry. For a physical disk with multiple paths, Disk Management will display each path as its own device, although each path is pointing to the same physical device.

For example, a dual-ported disk array connected to a single host may be seen by Disk Management as Disk 0 and Disk 1. When SVA Path is installed and these logical drives are not excluded, Disk 0 will be accessible while Disk 1 is inaccessible, as seen by Disk Management.

Note: A multiple-ported physical disk (such as a dual-ported fibre channel drive) is not a supported configuration within Microsoft Windows. Without SVA Path software, disk errors and data corruption will occur when data is written to the disk.

By understanding the relationship between the Windows' Registry, SVA Path's `setsp -a` output and Disk Administrator/Management, you will be able to track

system events back to a physical device, as will be explained further in Chapter 5., “Diagnosing Errors”, beginning on page 43.

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 36. 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 43.

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 drive. (pages 35 and 40). Must be used with -l -b0 disables load balancing for the device; -b1 enables it
-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 36 and 40)
-e{0 1}	set exclusion for a logical drive (pages 35 and 37). 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 36 and 40)
-g	generate configuration files after an <code>sppath</code> command (Page 32)
-i	show contents of driver configuration files

Table 4-1 *setsp Command Options*

Option	Effect
-L <parameter>	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 36)
-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, where <i>n</i> should be a number approximately equal to the device's maximum queue depth (pages 35 and 39). Must be used with -l.
-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 36 and 40)
-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	ignores spd devices (changes will take effect at next reboot). Ignored devices behave like drives not under the control of the SVA Path driver

The spmon Command

The `spmon` command is primarily associated with load monitoring and balancing of SVA Path functional devices (FDevs). FDev is a logical disk 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.

`spmon` is invoked automatically at boot time and runs as a Windows service. As such, events such as the reassignment of paths, are logged in the Application log. `spmon` can also be started or stopped via the Services icon in the Control Panel, or by issuing either `net stop spmon` or `net start spmon` from a Command Prompt.

Note: For optimum performance when striping data across multiple LUNs, disable load balancing on the devices being striped using the `setsp -b0` option.

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 or one half of the FDevs. If the path group includes one or more multi-FDev LUNs, each FDev is considered a separate device.

- `read-overhead μs -per-cmd μs -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 μs -per-cmd μs -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.

All device movements are logged to the system Event Logs, which can be accessed with the Event Viewer utility (see “Errors in the Event Viewer,” on page 43).

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 2 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 4-2 *Blocksize*

	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:

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

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/O's and higher if there are mostly large I/O's. 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    c3b0t1d0/6  Good    X           32    0       20    3000    1
      c4b0t0d0/12 Good
HardDisk 6  I:                               ID = "STK V960 0000000010390000"
=====
  1    c3b0t1d1/7  Excluded  X   X   32    0       20    3000    1
      c4b0t0d1/13 Excluded X   X
HardDisk 7                               ID = "STK V960 0000000010390001"
=====
  2    c3b0t1d2/8  Good    X           32    0       20    3000    1
      c4b0t0d2/14 Good    X
HardDisk 8                               ID = "STK V960 0000000010390002"
=====
```

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, the following command, `spmon show`, (Figure 4-2, “Example `spmon show` Output”) will show that each device in Figure 4-1, “Example `setsp -a` Output in Determining Load Balancing” is accessible via three paths, using HBAs "c2b0" and "c3b4":

```
C:\>spmon show
Total I/O load across monitored HBAs = 3ms

  HBA   Device   I/O Load  %load
-----
c2b0   c2b0         0ms    0.0%
c3b4   c3b4         3ms  100.0%

Device Harddisk1 [STK V960 0000000010850000]
FDev   Path           I/O Load  HBA Load
-----
spd0:0  1/Harddisk4     1ms       3ms  50.0%

Device Harddisk2 [STK V960 0000000010850001]
FDev   Path           I/O Load  HBA Load
-----
spd1:0  1/Harddisk5     1ms       3ms  50.0%

Device Harddisk3 [STK V960 0000000010850003]
FDev   Path           I/O Load  HBA Load
-----
spd2:0  2/Harddisk9     0ms       0ms  0.0%
```

Figure 4-2 Example `spmon show` Output

The column headings describe the various fields on the screen:

- **FDev** 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. For example, 350ms = 350 milliseconds

seconds (with an "s" suffix) if it is less than 300 seconds. For example, 18.1s = 28.1 seconds

hours:minutes:seconds (hh:mm:ss) if it is more than 59 minutes. For example, 0:12:42 = 0 hours, 12 minutes, and 42 seconds.

- **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, c3b4 is carrying 100% of the load. At the next measurement interval, spmon will automatically calculate and redistribute the I/O load among the three paths, significantly reducing the maximum load on c3b4.

Adding FDevs to an Existing Lun

After adding new FDevs to an existing lun via the "Add FDev" button on the SVA console, run the following command:

```
>spmon update-fdevs
```

This enables spmon to load balance the new FDevs separately from the rest of the FDevs that make up the LUN. If this command is not run, the FDevs will be treated as if they belonged to the last original FDev and will be prevented from being assigned to different paths. For example, if a LUN was originally composed of two FDevs spd0:0 and spd0:1 and the user adds three FDevs using the SVA console, the new FDevs are assigned to spd0:1.

The sspath Command

The sspath command is run automatically at boot after SVA Path is installed. sspath'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 {cXbYtZdN}	ignore the device or group of devices specified; devices are specified in the form <code>cX</code> , <code>cXbY</code> , <code>cXbYtZ</code> or <code>cXbYtZdN</code> . all devices that match will be ignored; Should be followed by <code>setsp -g</code> (see note below)
-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`. Changes will take effect at the next system reboot.

Display Options for `sppath`

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

```
C:\>sppath -v
SPD=0 c2b0t0d0 dev=0, 1 type=2 SANID="STK V960 0000000010850000"
SPD=0 c3b4t29d0 dev=0, 4 type=2 SANID="STK V960 0000000010850000"
SPD=0 c4b5t0d0 dev=0, 7 type=2 SANID="STK V960 0000000010850000"
SPD=1 c2b0t0d1 dev=0, 2 type=2 SANID="STK V960 0000000010850001"
SPD=1 c3b4t29d1 dev=0, 5 type=2 SANID="STK V960 0000000010850001"
SPD=1 c4b5t0d1 dev=0, 8 type=2 SANID="STK V960 0000000010850001"
SPD=2 c2b0t0d2 dev=0, 3 type=2 SANID="STK V960 0000000010850003"
SPD=2 c3b4t29d2 dev=0, 6 type=2 SANID="STK V960 0000000010850003"
SPD=2 c4b5t0d2 dev=0, 9 type=2 SANID="STK V960 0000000010850003"
```

Figure 4-3 *Using sppath in Verbose Mode*

Ignoring and Reclaiming Devices with sspath

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

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 <cXbYtZdN>
>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 reboot.

Understanding the setsp -a Screen

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

```
setsp -a
=====
      spd      Path/disk      Status Pri ExcBufBalanceRtrCntRtrDlyFailBack
=====
      0      c3b0t1d0/6      GoodX3202030001
          c4b0t0d0/12      Good
HardDisk 6  I:      ID = "STK V960 0000000010390000"
=====
      1      c3b0t1d1/7      ExcludedX3202030001
          c4b0t0d1/13      ExcludedXX
HardDisk 7      ID = "STK V960 0000000010390001"
=====
      2      c3b0t1d2/8      GoodX3202030001
          c4b0t0d2/14      GoodX
HardDisk 8      ID = "STK V960 0000000010390002"
=====
```

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 `HardDisk X`, where `X` is a number used by Disk Management, and the Drive Letter(s) assigned by Disk Management. This device number also appears in errors reported in the (System) Event Log.
- `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 both the disk number in the form `cWbXtYdZ` and the number assigned by **Disk Management**.

- **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.

Note: When Load-balancing is enabled all paths are active at all times, but there is no indication on this display.

- **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. For peak performance, **Buf** should be approximately equal to the maximum useful queue depth of the logical unit. Values between 1 and 100 are permitted. The default value is 16.
- **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.
- `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 28).

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 34, 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).

```
setsp -v -l0 -p1
=====
      spd      Path/disk      Status Pri  ExcBufBalanceRtrCntRtrDlyFailBack
=====
          0      c5b5t0d0/1    X3202030001
              c6b5t0d1/12
HardDisk 1              ID = "STK V9600 0000000010390000"
```

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 un-excluded) 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 refresh (or re-scan) disks in the Disk Management window to see the disk. When excluding a device, you must close and reopen the Computer Management window to update the Disk Management display.

“Unknown Devices” Errors in Windows 2003

Windows 2003 Disk Management may display excluded devices as “unknown.” Any attempt to initialize the “unknown” devices will fail, due to SVA Path device exclusion. To prevent this error, use the following steps to disable the “unknown” devices to Windows 2003.

1. Select **Computer Management > System Tools > Device Manager > Disk Drives.**
2. Right-click on the device you want to exclude, and click **Properties.**
3. Click the **General** tab.
4. Disable each device to be excluded.

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 36) and <n> is the new value. The system must be rebooted for this change to take effect. The number of buffer pointers recommended for a given device is approximately equal to the logical unit’s maximum useful queue depth; the default value of 16 should suit most SVA devices.

Removing LUNs from SVA Path

When a LUN on the SVA is deleted, SVA Path will continue to reference it until the next re-boot. If desired, use the

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

command to remove the reference immediately.

Excluding a Device Using the -x Option

The -x option causes a device to be inaccessible from SVA Path. The device then will still be under SVA Path's control but can only be accessed from the system. This option is useful with LUN masking.

1. To exclude a device using the -x option, run the

```
>setsp -x -l<spdN>
```

command, where spd# is the number of the SPD to be excluded.

2. Next, you need to stop the spd using the following command:

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

Note: Notice that the SPD is still in sspath.conf, but not in spd.conf, indicating that this spd is still under SVA Path’s control.

To check for devices that have been excluded using the -x option, type the following command:

```
>setsp -u2
```

To remove the device from the excluded devices list (displayed with the `u2` option), run the following commands:

```
>setsp -g -l<spdN>
```

```
>setsp -S -l<spdN>
```

Turning Load Balancing On and Off

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

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

where `<x>` 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<x> -r<n>
```

where `<x>` is the numerical element of the `spd` number (see “Specifying a Device for `setsp`” on page 36) and `<n>` 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<x> -d<n>
```

where `<x>` is the numerical element of the `spd` number and `<n>` 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<x> -f{0|1}
```

where `<x>` 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. Run the **Disk Management**. This will cause new devices to be detected by Windows.

Note: When Windows asks you to sign the disks, select No, and close the **Disk Management**.

2. Run this command to cause SVA Path to detect the new devices:

```
>sppath -v
```

3. To create an updated configuration file, run:

```
>setsp -g
```

4. To start the new device, run:

```
>setsp -S -l <new_spdN>
```

5. Un-exclude the new device:

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

6. Run `setsp -a` and check for a newly created spd number.

7. Run **Disk Management** again and sign each disk managed by SVA Path.

Deleting Devices

Note: Prior to deleting a device from SVA Path, you must use **Disk Management** to remove the device partition for the device to be deleted.

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(spdN)
```

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

```
#sppath -I cXbYtZdN
```

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

Errors in the Event Viewer

All changes in SVA Path devices and activity on devices under SVA Path's control are recorded as System events in the Windows Event Logs, accessed with the **Event Viewer** utility.

If a path to a storage system becomes unusable, the disk driver will report a "device not ready for access" error (Figure 5-6, "Sample disk driver error"), while the SPD driver will report an "SPD driver detects an IO path failure" warning (Figure 5-7, "Sample spd driver error").

An "SPD's TUR command failed" error will occur when the system is unable to communicate with the device after the system sends out a TUR (test unit ready) request. The error notes the name of the failed path, but does not pinpoint the source of the failure.

In the event of a path failure, use the **Disk Management** and **Disk Manager** utilities to troubleshoot the problem. A quick check of the bus on which the device had failed is also recommended.




Figure 5-6 Sample disk driver error




Figure 5-7 Sample spd driver error

VDS Errors in the System Event Log

The Virtual Disk Service (VDS) feature of Windows 2003 may issue an Event ID 6 error to the Windows Event Log when you start SVA Path. You can ignore these errors, as SVA Path does not use VDS.

Appendix A. Configuring SVA Path With Microsoft Cluster Service

This appendix provides recommendations for configuring SVA Path on a Windows 2000 server running Microsoft Cluster Service. Microsoft Cluster Service must be installed and configured prior to installing SVA Path.

Microsoft Cluster Service is included with Microsoft Windows 2000 Advanced Server. It provides failover support to the servers in the cluster, so that if one server fails, its workload is shifted to the other server.

Note: Although Microsoft Cluster Service and SVA Path are both available on Windows 2003 as well as Windows 2000, this appendix describes how to configure Microsoft Cluster Service and SVA Path on Windows 2000 only.

Note: This appendix does not describe Microsoft Cluster Service installation and configuration. For details on these topics, see the Microsoft Windows 2000 Advanced Server documentation.

Minimum System Requirements

Table A-1, “Minimum System Requirements”, on page 48 lists the hardware and software components required for running SVA Path with Microsoft Cluster Service.

Table A-1 Minimum System Requirements

Environment Item	Cluster Node 1	Cluster Node 2
Host operating system	Microsoft Windows 2000 Advanced Server, Version 5.00.2195 SP3, English version or Microsoft Windows Server 2003, English version	
Cluster management tool	Microsoft Cluster Service, Version 5.0	
SVA subsystem	V960, V2X/V2X2	
Virtual Power Suite software	For Windows 2000: SVA Path 3.3, Build 746 For Windows 2003: SVA Path 3.3, Build 748	

Environment Description

The Microsoft Cluster Service environment consists of two servers (referred to as nodes) and storage devices which are physically connected to the nodes. Figure A-1, “Sample Microsoft Cluster Service / SVA Subsystem Configuration”, on page 49 shows a sample clustering configuration involving two nodes and an SVA subsystem. Other configurations are possible.

Two types of network connections are used in this configuration:

- **Fibre connection**—Links each node to the SVA subsystem. This connection consists of a fibre wire that links the node HBA (host bus adapter) to the switch. There can be several HBAs on each node. This connection is referred to as the public network (not to be confused with the Ethernet public LAN).

- **Ethernet connection**—Links the nodes together. This connection consists of an Ethernet cable that links both nodes using their Ethernet cards. This network is used for clustering purposes. This connection is referred to as the private network or cluster network.

Each node of the cluster hosts the following components:

- SVA Path software
- Microsoft Cluster Service




Figure A-1 Sample Microsoft Cluster Service / SVA Subsystem Configuration

Configuring Microsoft Cluster Service

Hardware Configuration

Prior to installing SVA Path, there must be only one HBA per node connected to the SVA subsystem; therefore, you should disconnect any extra HBAs. If there is more than one HBA per node connected to the SVA subsystem, volumes will be duplicated. See Figure A-2, “Configuration Prior to Installing SVA Path”.




Figure A-2 Configuration Prior to Installing SVA Path

Preparing the Devices


Note: Both nodes in the cluster must have access to the same SVA domain.

Note: Microsoft Cluster Service supports basic volumes only.


Note: Microsoft Cluster Service uses a special disk—called the quorum disk—to store cluster information.

Use the Windows 2000 Disk Management console to discover and format the devices on both nodes.

1. From Node 1: Once you have created new devices, use the Disk Management console to have Windows 2000 detect them.



2. From Node 1: Create and format volumes on these devices.



3. From Node 2: Use the Disk Management console to have Windows 2000 detect the new volumes. Make the volume letters match the ones used by Node 1.

Disk 5 Basic 11.45 GB Online	Vol_B (Y:) 11.45 GB NTFS Healthy
Disk 6 Basic 11.45 GB Online	Vol_C (X:) 11.45 GB NTFS Healthy
Disk 7 Basic 11.45 GB Online	Vol_D (W:) 11.45 GB NTFS Healthy
Disk 8 Basic 2.28 GB Online	MSCS_Vol (V:) 2.28 GB NTFS Healthy

Microsoft Cluster Service Installation


During Microsoft Cluster Service installation you must select the volumes that will be shared by all the nodes of the cluster. It is recommended that you initially select only the quorum disk as a clustered resource. You will be able to add other volumes to the clustered resources later.

Note: For details on Microsoft Cluster Service installation, see the Microsoft Windows 2000 Advanced Server or Windows Server documentation.

Microsoft Cluster Service Default Resources

Once you have installed Microsoft Cluster Service on all the nodes, the minimum clustered resources are:

- Cluster IP address
- Cluster name
- Quorum disk



Note: For details on Microsoft Cluster Service resources, see the Microsoft Windows 2000 Advanced Server documentation.

Your cluster is now ready for SVA Path installation.

Installing and Configuring SVA Path

Note: When installing SVA Path on a node, it is highly recommended that you have only one HBA connected to the SVA subsystem.

Note: The term “active node” refers to the cluster node that owns disk resources.

1. From Node 1 (active node): Install the SVA Path software. Reboot the node once installation is complete.
2. From Node 2 (active node): Install the SVA Path software. Reboot the node once installation is complete.
3. From Node 1 (active node): Use SVA Path to configure the volumes:
 - a. Use the `sppath -v` command to verify that the new disks are acknowledged by SVA Path.

```
SPD=0 c4b1t0d0 dev=0, 3 type=2 SANID="STK V2X 7034005100000000"  
SPD=1 c4b1t0d1 dev=0, 4 type=2 SANID="STK V2X 7034009100000000"  
SPD=2 c4b1t0d2 dev=0, 5 type=2 SANID="STK V2X 7034003700000000"  
SPD=3 c4b1t0d3 dev=0, 6 type=2 SANID="STK V2X 7034007A00000000"  
SPD=4 c4b1t0d4 dev=0, 7 type=2 SANID="STK V2X 7034009200000000"  
SPD=5 c4b1t0d5 dev=0, 8 type=2 SANID="STK V2X 7034009300000000"
```

- b. Use the `setsp -g` command to update the SVA Path configuration for the new disks.
- c. Use the `setsp -S` command to start the driver for the new disks.
- d. Use the `setsp` command to display the volumes that can be used.

spd	Path/disk	Status	Pri	Exc	Buf	Balance	RtrCnt	RtrDly	FailBack
0	c4b1t0d0/3 HardDisk 3	Excluded	X	X	32	1	20	3000	1
						ID = "STK V2X 7034005100000000"			
1	c4b1t0d1/4 HardDisk 4	Excluded	X	X	32	1	20	3000	1
						ID = "STK V2X 7034009100000000"			
2	c4b1t0d2/5 HardDisk 5	Excluded	X	X	32	1	20	3000	1
						ID = "STK V2X 7034003700000000"			
3	c4b1t0d3/6 HardDisk 6	Excluded	X	X	32	1	20	3000	1
						ID = "STK V2X 7034007A00000000"			
4	c4b1t0d4/7 HardDisk 7	Excluded	X	X	32	1	20	3000	1
						ID = "STK V2X 7034009200000000"			
5	c4b1t0d5/8 HardDisk 8	Excluded	X	X	32	1	20	3000	1
						ID = "STK V2X 7034009300000000"			

- e. Use the `setsp e0 -lall` command to assign (un-exclude) the new disks.
- f. Use the `setsp` command to verify that the status of all paths is "Good."

spd	Path/disk	Status	Pri	Exc	Buf	Balance	RtrCnt	RtrDly	FailBack
0	c4b1t0d0/3 HardDisk 3	Good	X		32	1	20	3000	1
						ID = "STK V2X 7034005100000000"			
1	c4b1t0d1/4 HardDisk 4	Good	X		32	1	20	3000	1
						ID = "STK V2X 7034009100000000"			
2	c4b1t0d2/5 HardDisk 5	Good	X		32	1	20	3000	1
						ID = "STK V2X 7034003700000000"			
3	c4b1t0d3/6 HardDisk 6	Good	X		32	1	20	3000	1
						ID = "STK V2X 7034007A00000000"			
4	c4b1t0d4/7 HardDisk 7	Good	X		32	1	20	3000	1
						ID = "STK V2X 7034009200000000"			
5	c4b1t0d5/8 HardDisk 8	Good	X		32	1	20	3000	1
						ID = "STK V2X 7034009300000000"			

4. From Node 1 (active node): Reconnect all the node HBAs to the SVA subsystem.
5. From Node 1 (active node): Use SVA Path to reconfigure the volumes.

- a. Use the `sppath -v` command to verify that the new disks are acknowledged by SVA Path.


```
SPD=0 c4b1t0d0 dev=0, 3 type=2 SANID="STK V2X 7034005100000000"
SPD=0 c5b1t0d0 dev=0, 9 type=2 SANID="STK V2X 7034005100000000"
SPD=1 c4b1t0d1 dev=0, 4 type=2 SANID="STK V2X 7034009100000000"
SPD=1 c5b1t0d1 dev=0, 10 type=2 SANID="STK V2X 7034009100000000"
SPD=2 c4b1t0d2 dev=0, 5 type=2 SANID="STK V2X 7034003700000000"
SPD=2 c5b1t0d2 dev=0, 11 type=2 SANID="STK V2X 7034003700000000"
SPD=3 c4b1t0d3 dev=0, 6 type=2 SANID="STK V2X 7034007A00000000"
SPD=3 c5b1t0d3 dev=0, 12 type=2 SANID="STK V2X 7034007A00000000"
SPD=4 c4b1t0d4 dev=0, 7 type=2 SANID="STK V2X 7034009200000000"
SPD=4 c5b1t0d4 dev=0, 13 type=2 SANID="STK V2X 7034009200000000"
SPD=5 c4b1t0d5 dev=0, 8 type=2 SANID="STK V2X 7034009300000000"
SPD=5 c5b1t0d5 dev=0, 14 type=2 SANID="STK V2X 7034009300000000"
```

- b. Use the `setsp -g` command to update the SVA Path configuration for the new disks.
- c. Use the `setsp` command to verify that the status of all paths is "Good."

spd	Path/disk	Status	Pri	Exc	Buf	Balance	RtrCnt	RtrDly	FailBack
0	c4b1t0d0/3	Good	X		32	1	20	3000	1
	c5b1t0d0/9	Good			32	1	20	3000	1
HardDisk 3 ID = "STK V2X 7034005100000000"									
1	c4b1t0d1/4	Good	X		32	1	20	3000	1
	c5b1t0d1/10	Good			32	1	20	3000	1
HardDisk 4 ID = "STK V2X 7034009100000000"									
2	c4b1t0d2/5	Good	X		32	1	20	3000	1
	c5b1t0d2/11	Good			32	1	20	3000	1
HardDisk 5 ID = "STK V2X 7034003700000000"									
3	c4b1t0d3/6	Good	X		32	1	20	3000	1
	c5b1t0d3/12	Good			32	1	20	3000	1
HardDisk 6 ID = "STK V2X 7034007A00000000"									
4	c4b1t0d4/7	Good	X		32	1	20	3000	1
	c5b1t0d4/13	Good			32	1	20	3000	1
HardDisk 7 ID = "STK V2X 7034009200000000"									
5	c4b1t0d5/8	Good	X		32	1	20	3000	1
	c5b1t0d5/14	Good			32	1	20	3000	1
HardDisk 8 ID = "STK V2X 7034009300000000"									

6. Make Node 2 the active node.
7. From Node 2 (active node), perform Steps 3 and 4.
8. Verify that the volumes have the same letters on both nodes. If they do not, edit the drive letters to make them match.

9. You can now use the Windows 2000 Cluster Administrator console to define the volumes as clustered resources.



Using SVA Path in a Microsoft Cluster Service Environment

This section describes how to perform basic device management activities in the Microsoft Cluster Service environment. The following configuration is assumed:

- Nodes are clustered.
- SVA Path is installed and configured.
- There is more than one path on the nodes.

Adding New Volumes

1. From any host: Create devices using either the SVAA CLI (command-line interface) or the SVAC (Shared Virtual Array Console).
2. From Node 1 (active node): From the Windows 2000 Disk Management console, perform a Scan Hardware Changes operation. Each new disk should appear twice; therefore do not yet create volumes on the new disks.
3. From Node 1 (active node): Use SVA Path to configure the new disks:

- a. Use the `sppath -v` command to verify that the new disks are acknowledged by SVA Path.
 - b. Use the `setsp -g` command to update the SVA Path configuration for the new disks.
 - c. Use the `setsp -S` command to start the driver for the new disks.
 - d. Use the `setsp -e0` command to assign (un-exclude) the new disks.
4. From Node 1 (active node): Use the Windows 2000 Disk Management console to perform a Rescan Disks operation. Each new disk should appear only once.
 5. Turn Node 2 into the active node.
 6. From Node 2 (active node): Perform Steps 2, 3, and 4.
 7. From Node 2 (active node): Using the Windows 2000 Disk Management console, create and format volumes on the new disks. Make sure each volume has a letter.
 8. Turn Node 1 into the active node.
 9. From Node 1 (active node): Using the Windows 2000 Disk Management console, perform a Rescan Disks operation. The new volumes should be displayed. Verify that the volume letters on Node 1 match those on Node 2.
 10. Restart both nodes. (This is required by Microsoft Cluster Service. The volumes cannot be clustered without a reboot.) Once the nodes have been restarted, verify that all volume letters on Node 1 match those on Node 2.
 11. You can now use the Windows 2000 Cluster Administrator console to define the volumes as clustered resources.

Deleting Existing Volumes

Note: The active node is not relevant in this activity.

Note: Do not delete the quorum volume.

Note: After a disk has been deleted, its ID is still present in the SVA Path configuration. Consequently, a new volume may have the same ID as a deleted disk. This behavior is not specific to the Microsoft Cluster Service environment.

1. If the volumes to be deleted (target volumes) are clustered resources, use the Windows 2000 Cluster Administrator console to un-share them.
2. From Node 1: Use the Windows 2000 Disk Management console to perform a Rescan Disks operation. The target volumes should be displayed.
3. From Node 2: Perform Step 2.
4. From Node 1: Use the Windows 2000 Disk Management console to delete the target volumes.
5. From Node 1: Use the Windows 2000 Disk Management console to perform a Rescan Disks operation on both nodes. The target volumes should not appear.
6. From Node 1: Use the following SVA Path commands to stop the driver for the target disks.

```
setsp -T -l $nn$   
setsp -x -l $nn$ 
```

where nn is the ID of the SPD.

7. From Node 2: Perform Step 6.
8. From the server where SVAA is installed: Use either the SVAA CLI (command-line interface) or the SVAC (Shared Virtual Array Console) to delete the SVA devices that make up the target disks.
9. From Node 1: Use the Windows 2000 Disk Management console to perform a Scan Hardware Changes operation. A

Lost Device dialog box may appear. Target volumes should no longer appear in the list of managed devices.

10. From Node 2: Perform Step 9.



NEED MORE INFORMATION?
www.storagetek.com

ABOUT STORAGETEK

Storage Technology Corporation (NYSE: STK) is a \$2 billion global company that enables businesses, through its information lifecycle management strategy, to align the cost of storage with the value of information. The company's innovative storage solutions manage the complexity and growth of information, lower costs, improve efficiency and protect investments. For more information, visit www.storagetek.com, or call 1.800.275.4785 or 01.303.673.2800.

WORLD HEADQUARTERS

Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA
1.800.678.4430 or 01.303.673.4430

© 2004 Storage Technology Corporation, Louisville, CO. All rights reserved. Printed in USA. StorageTek and the StorageTek logo are registered trademarks of Storage Technology Corporation. Other names mentioned may be trademarks of Storage Technology Corporation or other vendors/manufacturers.

StorageTek equipment is manufactured from new parts, or new and used parts. In some cases, StorageTek equipment may not be new and may have been previously installed. Regardless, StorageTek's standard warranty terms apply, unless the equipment is specifically identified by StorageTek as "used" or "refurbished."

Replacement parts provided under warranty or any service offering may be either new or equivalent-to-new, at StorageTek's option. Specifications/features may change without notice.