Sun StorEdge™ Traffic Manager Installation and Configuration Guide

For the Solaris Operating Environment

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 816-1420-11
July 2003, Revision A

Submit comments about this document at: `http://www.sun.com/hwdocs/feedback`

Please
Recycle

Adobe PostScript

Contents

Preface v

1. Sun StorEdge Traffic Manager Overview and Architecture 1

What is the Sun StorEdge Traffic Manager Software? 2

The Multipath Solution 2

STMS Features 2

Multipath Configuration Management 3

I/O Load Balancing 3

Failover Support 3

Single Instance Multipath Devices 3

Dynamic Reconfiguration 3

STMS Architecture 4

Architecture Overview 4

STMS and the Solaris Operating Environment 5

Device Management Configurations 7

2. Installing the Sun StorEdge Traffic Manager Software 11

Downloading the Software 11

Downloading Packages 12

▼ To Download Packages 12

Downloading Patches 13

▼ To Download the Patches 13

Installing STMS 14

▼ To Check for Required Solaris 8 04/01 Patches 14

▼ To Install the Packages 14

▼ To Install the Patches 15

▼ To Verify the Installation. 15

3. Configuring the Sun Storage Traffic Manager Software 17

Configuring Ports Globally Through STMS 18

▼ To Enable Sun StorEdge Traffic Manger Globally 18

Configuring on a Per Port (pHCI) Basis 20

▼ To Enable or Disable a Port 20

▼ To Configure a Single PCI HBA 22

▼ To Configure a Dual PCI HBA 22

Configuring Third-Party Symmetric Storage Devices 24

▼ To Configure Third-Party Devices for Multipathing 24

▼ To Verify Third-Party Storage Devices Are Under STMS Control 26

Configuring Automatic Failback 27

▼ To Configure Automatic Failback Capability 27

▼ To Disable Automatic Failback Capability 28

Configuring Multipathing for Sun StorEdge T3 Array LUNs 29

▼ To Configure Multipathing Support for Sun StorEdge T3 LUNs 29

▼ To Configure Without Multipathing Support of Sun StorEdge T3 LUNs 30

Configuring Solstice DiskSuite or Solaris Volume Manager 31

Configuration Overview 31

▼ To Determine Pre-STMS Device Path Names 31

▼ To Unconfigure Solstice DiskSuite for Devices Under STMS Control 32

▼ To Enable STMS 32

▼ To Determine STMS Device Name to Pre-STMS Device Name Mapping 32

▼ To Reconfigure SDS 32

Example of Mirrored Devices 32

4. Deployment Scenarios 43

About Traffic Manager Configuration Samples 43

Configuration without Traffic Manager 44

Utility: format 44

Configuration with Traffic Manager 45

Utility: format 45

Utility: luxadm probe 46

Utility: luxadm display 46

Utility: luxadm failover 50

Example of Failover: 51

5. Troubleshooting STMS Software 55

STMS Not Running Properly 56

STMS Running Properly, but luxadm display and luxadm failover are Failing 56

Sun StorEdge T3 Arrays Do Not Show 57

System Failed During Boot with `scsi_vhci` Attachment 57

Group of A5x00s Connected, But They Still Appear Under Physical Paths in Format 58

System and Error Messages 58

Preface

The intended audience for this book is system administrators who are responsible for installing the server-side software, as well as adding and configuring new workstations.

Before You Read This Book

For late breaking information, see the

- *Sun StorEdge Traffic Manager Release Notes: For the Solaris Operating Environment* (817-0385)

How This Book Is Organized

This book contains the following sections:

Chapter 1 describes the architecture of the Sun StorEdge Traffic Manager software.

Chapter 2 describes how to install the Sun StorEdge Traffic Manager software.

Chapter 3 describes how to configure the Sun StorEdge Traffic Manager software.

Chapter 4 describes deployment scenarios.

Chapter 5 describes issues you may encounter that may require troubleshooting.

There is also a Glossary containing is a list of words and phrases and their definitions.

Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook2™ online documentation for the Solaris™ operating environment
- Other software documentation that you received with your system

Typographic Conventions

Typeface[1]	Meaning	Examples
`AaBbCc123`	The names of commands, files, and directories; on-screen computer output	Edit your `.login` file. Use `ls -a` to list all files. `% You have mail.`
`AaBbCc123`	What you type, when contrasted with on-screen computer output	`% su` `Password:`
AaBbCc123	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the *User's Guide.* These are called *class* options. You *must* be superuser to do this. To delete a file, type `rm` *filename.*

1 The settings on your browser might differ from these settings.

Shell Prompts

Shell	Prompt
C shell	*machine-name*%
C shell superuser	*machine-name*#
Bourne shell and Korn shell	$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number
Late Breaking	*Sun StorEdge Traffic Manager Release Notes*	817-0385

Accessing Sun Documentation

You can view, print, or purchase a broad selection of Sun documentation, including localized versions, at:

`http://www.sun.com/documentation`

Contacting Sun Technical Support

If you have technical questions about this product that are not answered in this document, go to:

`http://www.sun.com/service/contacting`

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

`http://www.sun.com/hwdocs/feedback`

Please include the title and part number of your document with your feedback:

Sun StorEdge Traffic Manager Installation and Configuration Guide: For the Solaris Operating Environment, part number 816-1420-11

CHAPTER 1

Sun StorEdge Traffic Manager Overview and Architecture

This chapter describes the architecture as well as the features of Sun StorEdge Traffic Manager software.

Topics in this chapter include:

- "What is the Sun StorEdge Traffic Manager Software?" on page 2
- "STMS Architecture" on page 4

What is the Sun StorEdge Traffic Manager Software?

The Sun StorEdge Traffic Manager software, a new addition to the kernel, is a client that manages devices that are accessed by a host through multiple paths.

This section contains the following topics:

- "The Multipath Solution" on page 2
- "STMS Features" on page 2

The Multipath Solution

Modern high performance I/O bus architectures are migrating from a host-centric model where storage is private to a single host towards the storage area network (SAN) model where storage is treated as a peer to host computers. In a SAN architecture, the storage is managed as a pool of resources which is shared among multiple heterogeneous hosts via a shared I/O interconnect. Combined with this is the new requirements to provide (reliability, availability, and serviceability) (RAS) performance metrics, as Solaris operating environment attempts to better support enterprise and high availability (HA) environments.

To increase both availability and I/O bandwidth to the storage pools, the Solaris operating environment platforms need to attach to these SANs by using multiple host controller interfaces (HCIs) and I/O interconnects controller interfaces. With the addition of STMS, however, the Solaris operating environment is now capable of supporting the multipath device configurations presented by these new SAN architectures.

STMS Features

The following are unique characteristics of the Sun StorEdge Traffic Manager software:

- Multipath configuration management
- I/O load balancing across a set of Host Bus Adapters (HBAs)
- Failover semantics in the Event of interconnect or controller failures
- Single instance multipath device support
- Dynamic Reconfiguration

Multipath Configuration Management

The STMS architecture enables mapping of available paths to client devices to be managed by the framework undertaken by the pHCI devices. The vHCI driver is also expected to supply interfaces for user-level system management applications to query and manage the pathset configurations being maintained by an instance of a vHCI.

I/O Load Balancing

The vHCI driver selects and routes I/O requests from its attached client devices to the pHCI device that provides the best transport services to the device. This routing decision needs to consider both the default pathset assigned to client device request as well as any routing policy such as round robin.

Failover Support

The vHCI and pHCI drivers are responsible for managing failovers, both automatic and manual. An example of an automatic failover is a cable that is unintentionally disconnected, and an example of a manual failover is an admin DR's system board containing a pHCI. Failover semantics need to be compatible with the support that is provided by the existing layered products. If an interconnect or device error is noted by a pHCI driver instance, the vHCI is notified of the loss of transport service provided by the pHCI.

Single Instance Multipath Devices

The principal feature provided by the vHCI nexus is a path-independent bus nexus under which multipath client devices can be attached. Client devices are created as they are registered with the framework by the pHCI devices.

This provides both a path-independent name in `/dev` and `/devices` directories as well as eliminating the need for layered drivers which consolidates multiple device instances into a pseudo-single instance device.

Dynamic Reconfiguration

STMS supports dynamic reconfiguration.

STMS Architecture

The Sun StorEdge Traffic Manager software implements a new set of interfaces for HCI drivers, known as the multipath driver interface (MDI). The new software provides a set of services for multipath configuration management.

This section includes the following topics:

- "Architecture Overview" on page 4
- "STMS and the Solaris Operating Environment" on page 5
- "Device Management Configurations" on page 7

Architecture Overview

The fundamental change in the STMS architecture is the reconstruction of the device tree to represent a multipath device (client) as a *single device instance* in the Solaris operating environment instead of *multiple instances* to represent all physical paths to the devices, as implemented in the current architecture. See FIGURE 1-1 and FIGURE 1-2 which illustrate the current and the new Solaris device tree representations.

virtual host controller interface (vHCI) drivers are pseudo nexus drivers which implement multipath services for a specific command set or bus architecture. A single instance of a vHCI driver exists for each command set which supports multipath devices.

vHCI drivers are provided with naming and transport services by one or more physical Host Controller Interface (pHCI) devices which share the common command set or bus architecture such as SCSI-3.

In the STMS architecture, an HBA is called pHCI. Also the state of a path is indicated as either ONLINE, OFFLINE, STANDBY. The paths are now classified in terms of preference.

The principal feature provided by the vHCI nexus is a path independent bus nexus under which multipath client devices can be attached. Client devices are created as they are registered with the framework by the pHCI devices. This provides a path independent name in both `/dev` and `/devices` directories.

STMS delivers an implementation of a vHCI driver called scsi_vhci for SCSI-3 Fibre Channel devices. The name of the node in the Solaris device tree is:

```
/devices/scsi_vhci
```

with client devices (targets) having names of the form:

```
/devices/scsi_vhci/ssd@g200000203709C3F5:a
```

STMS and the Solaris Operating Environment

FIGURE 1-1 and FIGURE 1-2 illustrate the current and the new Solaris device trees respectively. The illustration shows a system with dual-pathed Fibre Channel A5x00 storage arrays.

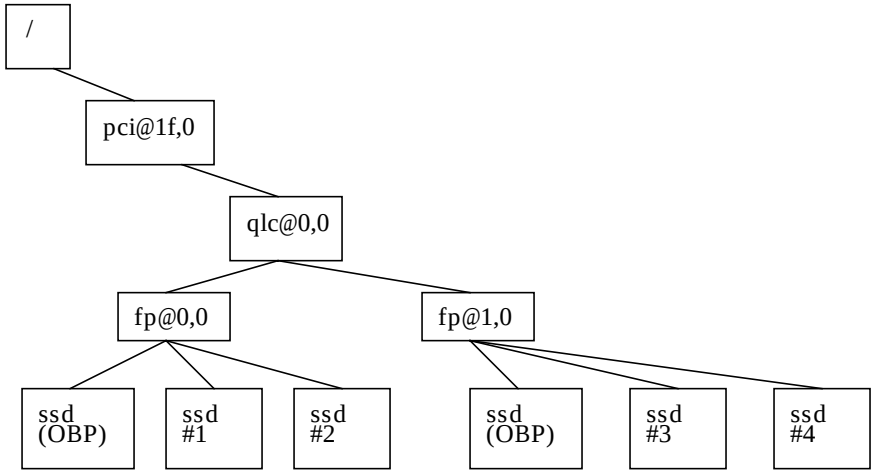

FIGURE 1-1 Current Device Tree With a Multipath Device

```
ssd#1 -> ssd@w210000203709C3f5,0
ssd#2 -> ssd@w210000203709D3f5,0
ssd#3 -> ssd@w220000203709C3f5,0
ssd#4 -> ssd@w220000203709D3f5,0
```

Note – The presence of two instances of ssd target device for a single dual-pathed Fibre Channel Device.

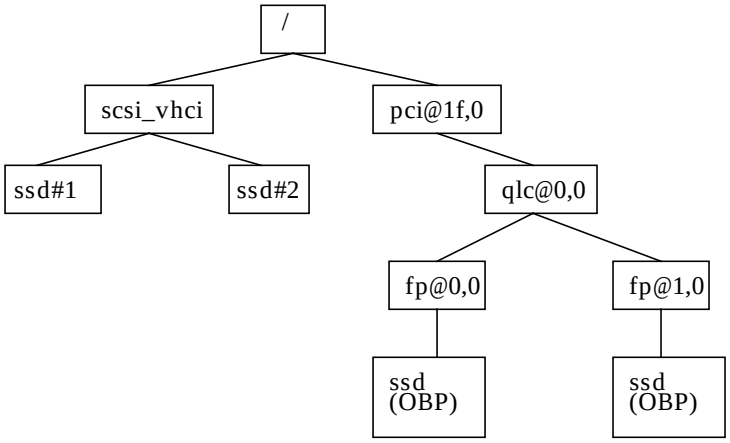

FIGURE 1-2 Subsection of Solaris Device Tree Under Sun StorEdge Traffic Manager

```
ssd#1 -> ssd@g200000203709C3f5,0
ssd#2 -> ssd@g200000203709D3f5,0
```

Note – ssd#1 and ssd#2 from FIGURE 1-1 have been consolidated to one ssd node ssd#1 under scsi_vhci similarly ssd#3, and ssd#4 have been consolidated to ssd#2 under scsi_vhci.

Device Management Configurations

FIGURE 1-3 and FIGURE 1-4 show two device management systems: one incorporates Sun StorEdge Traffic Manager software, whereas the other does not.

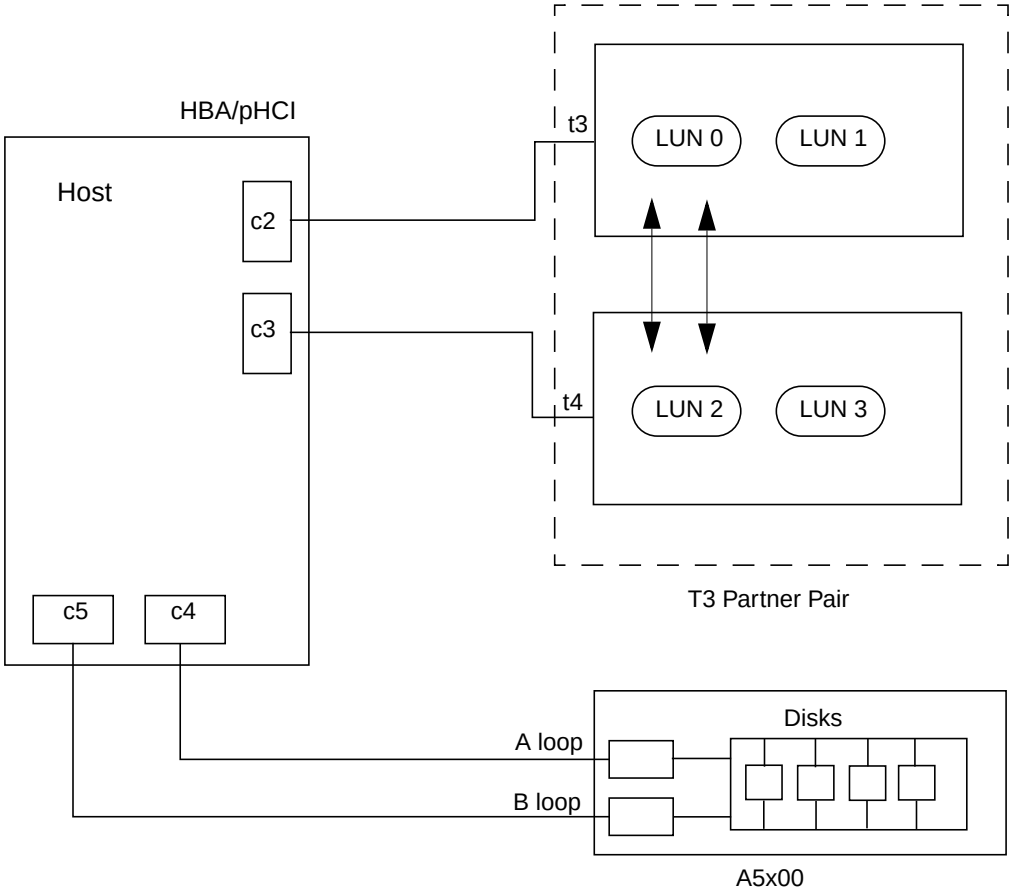

FIGURE 1-3 Device Management Configuration Without Traffic Manager

FIGURE 1-3 shows multiple paths that connect devices to their corresponding host.

In this configuration, the format of the path from the host to a device node is "c*x*t*y*d*z*," where *x* is a unique *controller number*, *y* is the *target id* that contacts the device, and *z* is the *LUN number* associated with that device.

For example, the path to refer the host to LUN 0 device node is "c2t3d0."

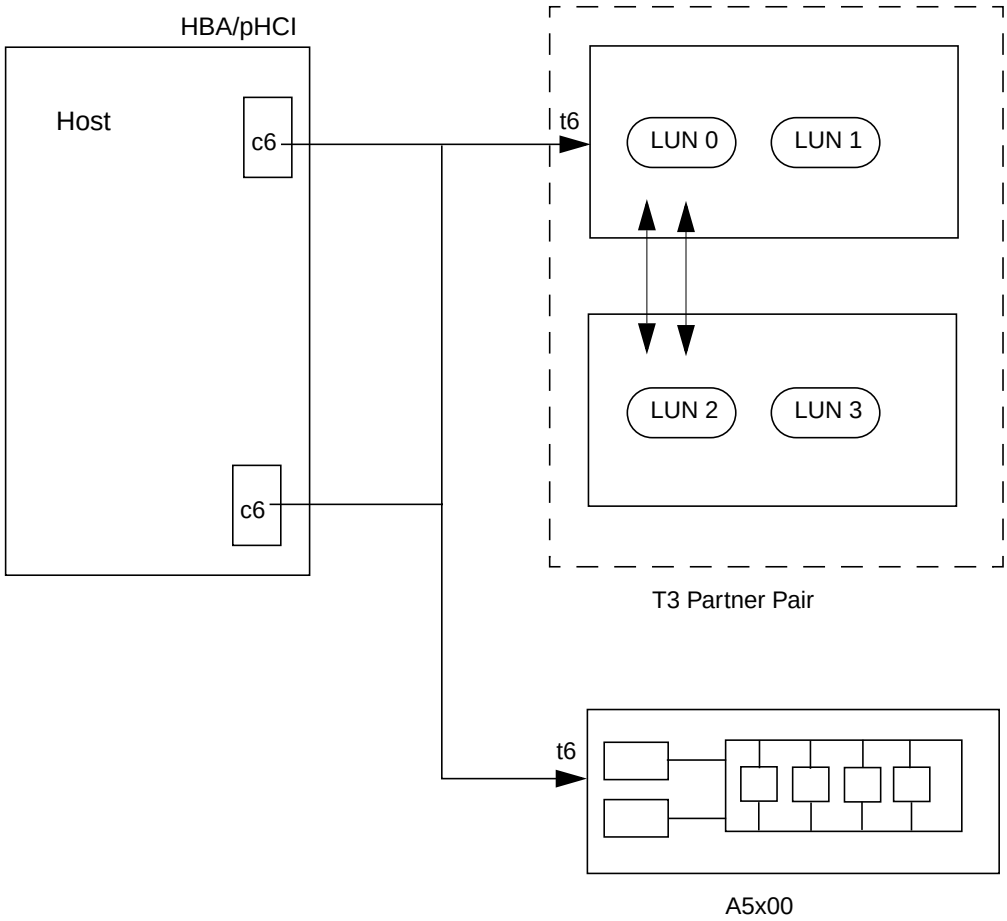

FIGURE 1-4 Device Management Configuration *With* Traffic Manger

Contrary to the previous configuration, FIGURE 1-4 shows that only one virtual controller exist for all devices that are managed by STMS. The virtual path entry hides the complexity of multiple paths to the same device.

In this configuration, the format of the path from the host to a device node is "c*x*tGUIDd*y*," where *x* is a *virtual controller number*, GUID is the Global Unique Identifier for the device, and *z* is the *LUN number* associated with that device.

A GUID is globally unique for each lun. GUID for the device is constructed from the INQUIRY VPD Page 0x83 data returned by the RAID array for that lun device. For example, the device names of the 2 luns of a Sun StorEdge T3 will have a long name as follows:

```
c6t60020F20000002E63EAFF9A000022560d0
c6t60020F20000002E63EAFF9BB000B8A16d0
```

For Sun StorEdge T3 LUNs, GUID is a long name that uniquely identifies each Sun LUN. For example, the path to connect the host to T3 LUN 0 device node is “c6t600...42d0.”

For JBODs, GUID is of the form World Wide Name. For example, the path to connect the host to A5x00 device with WWN of 200000020371A1862 is "c6t200000020371A1862d0".

In the configuration shown in FIGURE 1-4, exactly one virtual path entry exists for each device. The multiple paths, however, are not displayed to simplify the device tree.

You may use the `luxadm display` command for details about a device or LUN, as well as its associated multiple paths and status.

CHAPTER 2

Installing the Sun StorEdge Traffic Manager Software

This chapter contains instructions for a new installation of the Sun StorEdge Traffic Manager software. Before installing the software, ensure that your system meets the requirements identified in the *Sun StorEdge Traffic Manager Release Notes: For the Solaris Operating Environment*.

This section contains the following topics:

- "Downloading the Software" on page 11
- "Installing STMS" on page 14

Downloading the Software

The installation includes three packages and associated man pages. The following are the included packages:

- Fibre Channel stack
- STMS
- Fibre Channel HBA driver
- luxadm utility

This section contains the following topics:

- "Downloading Packages" on page 12
- "Downloading Patches" on page 13

Downloading Packages

TABLE 1 lists the packages and their location. You need to install the packages before you install the patches. If the packages are already on the system from a previous installation, you do not need to install them again.

Check `http://www.sun.com/download/` or `http://www.sun.com/storage/san` for updates. There is one set of packages for the Solaris 8 operating environment and another for the Solaris 9 operating environment available under the respective links for the operating environments. The `SUNWsan` package is interchangeable between the releases.

TABLE 1 Packages

Packages	Location

`SUNWsan`
`SUNWcfpl`
`SUNWcfplx`
	`http://www.sun.com/download/`
	or
	`http://www.sun.com/storage/san`

▼ To Download Packages

1. **Go to `http://www.sun.com/download/`.**

   The browser is redirected to `http://www.sun.com/software/download`.

2. **Under Browse Downloads, click View All.**
3. **Click SAN 4.*x*.**
4. **Register, if you have not already; otherwise enter your registration information.**
5. **Click Accept and Continue on the Legal/License Agreement page.**
6. **Click Download All.**

Downloading Patches

TABLE 2 contains a list of the minimum level patches necessary to use Sun StorEdge Traffic Manager software.

Check `http://sunsolve.sun.com` for updates.

TABLE 2 Patches

Name	Solaris operating environment 8 Patches	Solaris operating environment 9 Patches
Sun StorEdge Traffic Manager patch	111412-11	113039-03
`fctl/fp/fcp/usoc` driver	111095-13	113040-04
fcip driver	111096-06	113041-03
qlc driver	111097-12	113042-04
`luxadm/liba5k` and `libg_fc` patch	111413-10	113043-04
`cfgadm` fp plug-in library patch	111846-06	113044-03
SAN Foundation Kit patch	111847-07	111847-07

▼ To Download the Patches

1. **Go to** `http://sunsolve.sun.com`**.**
2. **Download the patches listed in** TABLE 2**.**

   You can access these by navigating the links PatchPro -> Network Storage Products. Check the box next to the adapter of your choice and then click Generate Patch List.
3. **Unzip the patches in the** `/tmp` **directory.**

Installing STMS

Installing STMS consists of checking for required patches, installing STMS packages and patches, and verifying the installation.

- "To Check for Required Solaris 8 04/01 Patches" on page 14
- "To Check for Required Solaris 8 04/01 Patches" on page 14
- "To Install the Patches" on page 15
- "To Verify the Installation." on page 15

▼ To Check for Required Solaris 8 04/01 Patches

1. **Verify installation of patches required Solaris 8 4/01.**

```
# patchadd -p | nawk '{print $2}' | grep 108982

# patchadd -p | nawk '{print $2}' | grep 108984

# patchadd -p | nawk '{print $2}' | grep 109529
```

   The minimum patch levels for these patches are: 108982-09, 108984-08, 109529-06.

2. **If the above patches are not present, download and install them before proceeding with installing STMS.**
   - These patches are available at `http://www.sun.com/sunsolve`.
   - Please study the README of these patches.

▼ To Install the Packages

- **Use the** `pkgadd` **command to install the packages in** TABLE 1**.**

  Read the README file for the packages and follow the instructions.

```
# pkgadd -d SUNWsan
# pkgadd -d SUNWcfpl
# pkgadd -d SUNWcfplx
```

▼ To Install the Patches

Install the packages before you install the patches.

1. **Use the** `patchadd` **command to install the patches for the operating environment you are using in** TABLE 2**.**

   You must install the patches in the order in which they are listed in TABLE 2.

2. **Expand the downloaded files.**

   If the filename suffix is `.tar.Z` then type:

```
# uncompress filename.Z
# tar -xvf filename.tar
```

   If the filename postfix is `.zip` then enter:

```
# unzip filename.zip
```

   The STMS patches are located in the directory */download_directory*.

3. **Add the five patches to the system in the following order**:

```
# patchadd /download_directory/111412-11
# patchadd /download_directory/111095-13
# patchadd /download_directory/111096-06
# patchadd /download_directory/111097-12
# patchadd /download_directory/111413-10
```

▼ To Verify the Installation.

1. **Check the `/var/sadm/patch` logs to verify installation has no errors.**
2. **Check for the existence of a scsi_vhci entry in the `name_to_major` file.**

```
# grep scsi_vhci /etc/name_to_major
```

   If an entry does not exist and no errors in patch logs, then call service.

You have completed installing the Sun StorEdge Traffic Manager software. Proceed to Chapter, “Post Installation Tasks.” If, however, you would like to leave your STMS system in the default disabled state, or if do not wish to configure your system, you may then reboot the system.

CHAPTER 3

Configuring the Sun Storage Traffic Manager Software

Configuring STMS depends on how you intend to use your system. STMS can be configured to control all of your Sun supported Fibre Channel HBAs as listed in the *Sun StorEdge Traffic Manager Release Notes: For the Solaris Operating Environment*

STMS can also be configured to control a subset of your Sun supported HBAs and is capable of performing load balancing over the Sun supported HBAs that it controls. The global configuration file contains a switch to either turn on or off load balancing. This section describes what files and parameters are used to configure STMS.

You can now use STMS to configure multipathing for third-party symmetric storage devices in addition to Sun storage devices.

This chapter also provides the instructions for configuring your Sun StorEdge Traffic Manager software if your system is using the Solstice DiskSuite or Solaris Volume Manager or is using a volume manager that is not supported by your STMS system.

Topics in this chapter include:

- "Configuring Ports Globally Through STMS" on page 18
- "Configuring on a Per Port (pHCI) Basis" on page 20
- "Configuring Third-Party Symmetric Storage Devices" on page 24
- "Configuring Automatic Failback" on page 27
- "Configuring Multipathing for Sun StorEdge T3 Array LUNs" on page 29
- "Configuring Solstice DiskSuite or Solaris Volume Manager" on page 31

Note – In a non-fabric environment, whenever "`mpxio-disable`" field in `scsi_vhci.conf`, `qlc.conf` or T3's `mp-support` field is changed, the host must go through a reconfiguration reboot.

Configuring Ports Globally Through STMS

Unless explicitly enabled or disabled, the global port settings in STMS will apply. You can enable STMS to manage ports through global settings and over-ride those settings with specific controls. See “Configuring on a Per Port (pHCI) Basis” on page 20 for more information.

▼ To Enable Sun StorEdge Traffic Manger Globally

1. **Using any text editor, edit the** `/kernel/drv/scsi_vhci.conf` **file:**

```
#
# Copyright (c) 2001 by Sun Microsystems, Inc.
# All rights reserved.
#
#pragma ident   "@(#)scsi_vhci.conf     1.3     01/08/22 SMI"
#
name="scsi_vhci" class="root";
#
# mpxio global enable/disable switch: setting mpxio-disable="no" will activate
# I/O multipathing; setting mpxio-disable="yes" disables this feature (do
# not remove this property).
#
mpxio-disable="yes";
#
# Load balancing global configuration: setting load-balance="none" will cause
# all I/O to a given device (which supports multipath I/O) to occur via one
# path.  Setting load-balance="round-robin" will cause each path to the device
# to be used in turn.
#
load-balance="round-robin";
#
# Force load driver to support hotplug activity (do not remove this property).
#
ddi-forceattach=1;
```

Do not change the name and class definitions.

2. **To enable STMS support for third party symmetric devices, make an entry like this:**

```
"SUN     StorEdge 3510", "symmetric-option";
symmetric-option = 0x1000000;
```

3. **To enable STMS globally, change the value of “mpxio-disable” to “no”.**

   By default “`mpxio-disable`” is set to “yes” which means STMS is disabled.

4. **If you want STMS to use all the available paths for load balancing, leave load-balance set to the default of “round-robin;” otherwise, change the definition to “none.”**

5. **Reboot Your System**

   After completing the procedures in this section you must reboot the system.

   In non-fabric environments, whenever "mpxio-disable" field in `scsi_vhci.conf` or T3s’ mp-support field is changed, the host must go through a reconfiguration reboot.

```
ok> boot -r
```

Configuring on a Per Port (pHCI) Basis

Even though the Sun StorEdge Traffic Manager software is globally enabled, per port STMS settings take precedence. Sun StorEdge Traffic Manager software will not take control of a port if it is explicitly disabled. A port that is neither explicitly enabled or disabled will be governed by the global the Sun StorEdge Traffic Manager software configuration. For the Sun StorEdge Traffic Manager software to take affect, it must be first globally enabled.

Load balancing is controlled by the global variable and cannot be applied on a per port basis.

Volume managers that do not support the Sun StorEdge Traffic Manager software must have STMS disabled on their HBA ports. This section contains the following procedures:

- "To Enable or Disable a Port" on page 20
- "To Configure a Single PCI HBA" on page 22
- "To Configure a Dual PCI HBA" on page 22

▼ To Enable or Disable a Port

1. **Log in as superuser.**

   Determine the HBAs to which you want your Sun StorEdge Traffic Manager software control.

2. **Open the `qlc.conf` configuration file in a text editor.**

```
# vi /kernel/drv/qlc.conf
```

3. **Explicitly Enable or Disable an HBA.**

   Add the property "mpxio-disable" to the HBA definition. To enable STMS on the port, set "mpxio-disable" to "no". To disable STMS on the port, set "mpxio-disable" to "yes."

For example, to select the correct device, perform a `ls -l` on `/dev/fc`. The following is an example of the output.

```
lrwxrwxrwx   1 root     root          49 Apr 17 18:14 fp0 ->
../../devices/pci@6,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx   1 root     root          49 Apr 17 18:14 fp1 ->
../../devices/pci@7,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx   1 root     root          49 Apr 17 18:14 fp0 ->
../../devices/pci@6,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx   1 root     root          49 Apr 17 18:14 fp1 ->
../../devices/pci@7,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx   1 root     root          49 Apr 17 18:14 fp2 ->
../../devices/pci@a,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx   1 root     root          49 Apr 17 18:14 fp3 ->
../../devices/pci@b,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx   1 root     root          50 Apr 17 18:14 fp4 ->
../../devices/pci@12,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx   1 root     root          56 Apr 17 18:14 fp5 ->
../../devices/pci@13,2000/pci@2/SUNW,qlc@4/fp@0,0:devctl
lrwxrwxrwx   1 root     root          56 Apr 17 18:14 fp6 ->
../../devices/pci@13,2000/pci@2/SUNW,qlc@5/fp@0,0:devctl
```

Note – The last two entries have two `/pci` elements. This is a dual PCI HBA. The rest of the entries do not have additional PCI bridges and is are Single PCI HBAs.

---

4. Reboot Your System

After completing the procedures in this section you must reboot the system.

In non-fabric environments, whenever "mpxio-disable" field in `scsi_vhci.conf` or T3s' mp-support field is changed, the host must go through a reconfiguration reboot.

```
ok> boot -r
```

▼ To Configure a Single PCI HBA

To configure the first entry from the list above, (that is, fp0 -> ../../devices/pci@6,2000/SUNW,qlc@2/fp@0,0:devctl) which is a single PCI HBA, add the following line to the `qlc.conf` file.

1. **To explicitly enable STMS on this port, enter:**

```
name="qlc" parent="/pci@6,2000" unit-address="2" mpxio-disable="no"
```

2. **To explicitly disable STMS on this port, enter:**

```
name="qlc" parent="/pci@6,2000" unit-address="2" mpxio-disable="yes"
```

3. **Reboot Your System**

After completing the procedures in this section you must reboot the system.

In non-fabric environments, whenever "mpxio-disable" field in `scsi_vhci.conf` or T3s' mp-support field is changed, the host must go through a reconfiguration reboot.

```
ok> boot -r
```

▼ To Configure a Dual PCI HBA

To configure the last entry from the above list, (`fp6 -> ../../devices/pci@13,2000/pci@2/SUNW,qlc@5/fp@0,0:devctl`) which is a dual PCI HBA, add the following line to the `qlc.conf` file.

1. **To explicitly enable STMS on this port, type:**

```
name="qlc" parent="/pci@13,2000/pci@2" unit-address="5" mpxio-disable="no”
```

2. **To explicitly disable STMS on this port, type:**

```
name="qlc" parent="/pci@13,2000/pci@2" unit-address="5" mpxio-disable="yes”
```

Devices with STMS enabled are enumerated under `/devices/scsi_vhci`. Those with STMS disabled are enumerated under physical path names.

Caution – All paths to storage devices either have STMS enabled or disabled.

Configuring STMS on a port basis enables STMS to coexist with other multipathing solutions like Alternate Pathing (AP). However, storage devices and paths should not be shared between STMS and other multipathing solutions.

Caution – Do not intermix any other multipathing software with STMS to the same storage device. Partial configuration of some paths to a storage device with STMS and other paths to the same storage device with other multipathing software is prone to cause system problems.

3. Reboot the system.

In a non-fabric environment, whenever "`mpxio-disable`" field in `scsi_vhci.conf`, `qlc.conf` or T3's `mp-support` field is changed, the host must go through a reconfiguration reboot.

```
ok> boot -r
```

Configuring Third-Party Symmetric Storage Devices

You can now use STMS to provide multipathing and load-balancing with more third-party symmetric storage devices than in earlier releases, enabling greater device heterogeneity in your SAN or direct-connect storage environment. Formerly, STMS worked only with the Sun StorEdge 5200, 3500 FC, T3 and T3+ arrays and the Sun StorEdge 39*x*0, 69*x*0, and 99**x**0 series. Now, you can use STMS with other symmetric storage devices, too.

To use the new functionality, you must edit new parameters in the `scsi_vhci.conf` file. This section contains the following procedures:

- "To Configure Third-Party Devices for Multipathing" on page 24
- "To Verify Third-Party Storage Devices Are Under STMS Control" on page 26

▼ To Configure Third-Party Devices for Multipathing

1. **Open the `scsi_vhci.conf` file in your text editor.**
2. **Enable `mpxio` and set load balancing by changing the following setting from:**

```
mpxio-disable="yes";
```

to:

```
mpxio-disable="no";
```

3. **Add the** `vendor_id` **and** `product_id` **properties.**

   The vendor ID (*v_id*) must be five characters and the product ID (*prod_id*) must be eight characters with trailing blanks:

```
device-type-scsi-options-list =
"v_id   prod_id ", symmetric-option",
symmetric-option = 0x1000000
```

   You can obtain the values for the storage *vendor_id* and *product_id* variables with the Inquiry command on your storage system. You must replace the variables with appropriate values for your system. For example:

```
device-type-scsi-options-list =
"ven-a SENA    ", "symmetric-option"
"ven-b SYMMET  ", "symmetric-option"
"ven-c OPEN    ", "symmetric-option"
symmetric-option = 0x1000000
```

4. **Save and close the** `scsi_vhci.conf` **file.**

5. **Reboot Your System**

   After completing the procedures in this section you must reboot the system.

   In non-fabric environments, whenever "mpxio-disable" field in `scsi_vhci.conf` or T3s' mp-support field is changed, the host must go through a reconfiguration reboot.

```
# boot -r
```

▼ To Verify Third-Party Storage Devices Are Under STMS Control

1. **After you configure the** `scsi_vhci.conf` **file, confirm that the devices nodes are created under the** `/devices/scsi_vhci` **file:**

```
# format
```

2. **If the devices are not created, confirm the** *vendor_id* **and** *product_id* **fields in the** `scsi_vhci.conf` **file.**

3. **Reboot Your System**

   After completing the procedures in this section you must reboot the system.

   In non-fabric environments, whenever "mpxio-disable" field in `scsi_vhci.conf` or T3s' mp-support field is changed, the host must go through a reconfiguration reboot.

```
# boot -r
```

Configuring Automatic Failback

In the earlier release of STMS, you had to manually failback a primary path with the `luxadm failover primary` *disk_name* command. Now you can configure the `scsi_vhci.conf` file so that primary paths fail back automatically when they become available after being in an OFFLINE state for a while. This feature is available on Sun StorEdge T3+ arrays.

- "To Configure Automatic Failback Capability" on page 27
- "To Disable Automatic Failback Capability" on page 28

▼ To Configure Automatic Failback Capability

1. **Open the `scsi_vhci.conf` file in your text editor.**
2. **Enable the automatic failback property.**

   The default setting is `"disable"`.

   ```
   auto-failback = "enable"
   ```

3. **Save and close the `scsi_vhci.conf` file.**

   When automatic failback is enabled, the following message is printed in the `/var/adm/messages` file during boot-up:

   ```
   /scsi_VHCI (scsi_vhci0): Auto-failback capability enabled through
   scsi_vhci.conf file.
   ```

4. **Reboot the system so that the changes will take effect:**

   ```
   # boot -r
   ```

▼ To Disable Automatic Failback Capability

1. **Open the** `scsi_vhci.conf` **file in your text editor.**
2. **Disable the automatic failback property:**

```
auto-failback = “disable”
```

3. **Save and close the** `scsi_vhci.conf` **file.**
4. **Reboot the system so that the changes will take effect:**

```
# reboot -- -r
```

Configuring Multipathing for Sun StorEdge T3 Array LUNs

For T3 LUNs to be enumerated under STMS, you need to set `mp_support` to "mpxio." This setting is only valid when the host has STMS enabled.

- "To Configure Multipathing Support for Sun StorEdge T3 LUNs" on page 29
- "To Configure Without Multipathing Support of Sun StorEdge T3 LUNs" on page 30

▼ To Configure Multipathing Support for Sun StorEdge T3 LUNs

- **On the Telnet or console of Sun StorEdge T3 array, type `sys list` to verify mp_support field. If it is not set to "mpxio," type the following command:**

```
t31> sys mp_support mpxio
```

Note – If `mp_support` is set to anything other than "`mpxio`" (like "rw" or "none"), Sun StorEdge T3 Array LUNs are enumerated in legacy non-STMS mode.

In this mode, LUN names will look longer under STMS. The long names look something like g60020f20000002253b22101d00029fc9, where each is called GUID (Global Unique Identifier). The `format` or `luxadm` commands will display T3 LUNs with this GUID under `scsi_vhci`.

For example, the `format` command shows the STMS T3 LUN as:

```
4. c4t60020F20000002253B22101D00029FC9d0 <SUN-T300-0117 cyl 34145
alt 2 hd 32 sec 128>
/scsi_vhci/ssd@g60020f20000002253b22101d00029fc9
```

Note – Set T3's `mp_support` to "`mpxio`," only when the host has STMS enabled. If host has STMS disabled, any I/O to a Sun StorEdge T3 array in "`mpxio`" mode will fail until `mp_support` is set to a desired mode (like rw or none).

▼ To Configure Without Multipathing Support of Sun StorEdge T3 LUNs

1. **On the Telnet or console of Sun StorEdge T3 array, type** `sys list` **to verify** `mp_support` **field.**

2. **If it is set to "mpxio," type the following command where** *any-value* **is anything but "mpxio":**

```
t31> sys mp_support any-value
```

If Sun StorEdge T3 LUNs do not need to be under STMS, the `mp_support` field should be set to anything other than "mpxio," and will not have any GUID associated with T3s.

They will be enumerated as legacy devices.

For example, if mp_support is rw, T3 LUNs are seen as:

```
4.     c2t1d3 <SUN-T300-0117 cyl 34145 alt 2 hd 32 sec 128>

/pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0/ssd@w50020f2300000225,3
8.     c3t2d3 <SUN-T300-0117 cyl 34145 alt 2 hd 32 sec 128>
       /pci@1f,2000/SUNW,qlc@1/fp@0,0/ssd@w50020f23000001f6,3
```

3. **Reboot**

In non-fabric environment, whenever Sun StorEdge T3's `mp_support` variable is changed to either enable or disable STMS, a *reconfiguration reboot* is needed on the host.

Configuring Solstice DiskSuite or Solaris Volume Manager

Solstice DiskSuite can be re configured to allow the Sun StorEdge Traffic Manager software control the Sun supported HBAs it uses. The disk usage information you saved prior to enabling STMS will be used. Using the `luxadm` command you can determine the mapping from the new STMS device path name to the pre-STMS device path name. This mapping will be used to reconfigure Solstice DiskSuite to recognize the correct devices under the new STMS device path names.

For more information concerning Solstice DiskSuite check the Solstice DiskSuite online manuals at:

`http://docs.sun.com`

This section contains the following topics:

- "Configuration Overview" on page 31
- "Example of Mirrored Devices" on page 32

Configuration Overview

The following are the steps to complete the configuration process:

Note – Backing up your data is recommended.

- "To Determine Pre-STMS Device Path Names" on page 31
- "To Unconfigure Solstice DiskSuite for Devices Under STMS Control" on page 32
- "To Enable STMS" on page 32
- "To Determine STMS Device Name to Pre-STMS Device Name Mapping" on page 32
- "To Reconfigure SDS" on page 32

▼ To Determine Pre-STMS Device Path Names

1. **Use the `metadb` command to get a list of devices used to store the Solstice DiskSuite configuration.**
2. **Use the `metastat -p` command to get the disk to meta device mapping.**

▼ To Unconfigure Solstice DiskSuite for Devices Under STMS Control

1. **Unmount the Solstice DiskSuite metadevices that will be under control of STMS.**
2. **Clear the metadevices using the `metaclear` command.**
3. **Clear the metadevice database using the `metadb -df` command which takes the list of disks, from the `metadb` command in Step 1 as a list of arguments.**

▼ To Enable STMS

- **See "Configuring Ports Globally Through STMS" on page 18.**

▼ To Determine STMS Device Name to Pre-STMS Device Name Mapping

- **Use the `luxadm display` command to determine which pre-STMS device names are combined under each STMS device name.**

▼ To Reconfigure SDS

- **Create metadevices using STMS path names with the `metainit` command.**

Caution – For RAID5 devices, be sure to used the "-k" option to prevent initializing the disks.

Example of Mirrored Devices

The following example shows how to migrate a two-way mirrored meta-devices from non-STMS to STMS.

A Sun StorEdge T3 partner pair is connected to the host, and mp_support on Sun StorEdge T3 is set to rw. STMS is initially disabled.

Four LUNs of equal size are created on the partner pair. Two metadb replicas are created on lun 0 (c2t1d0) and lun 1 (c2t1d1).

d10 and d11 are the sub-mirror meta devices created on the lun 2 (c2t1d2) and Lun 3 (c2t1d3).

d14 is the mirror of d10 and d11.

1. **Check paths.**

```
% cat /etc/lvm/md.tab
        d10             1 2 /dev/dsk/c2t1d2s1 /dev/dsk/c2t1d2s6
        d11             1 2 /dev/dsk/c2t1d3s1 /dev/dsk/c2t1d3s6
        d14             -m d10 d11
```

2. **Setup metadatabase.**

```
% metadb -a -f -c 2 /dev/dsk/c2t1d0s7 /dev/dsk/c2t1d1s7
```

3. **Set up the submirrors and mirror.**

```
% metainit d10
% metainit d11
% metainit d14
```

4. **Save the pre-STMS device information.**

Collect the output of `format`, `metadb`, `metastat`, and `metastat -p` commands.

```
% format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
          /pci@1f,4000/scsi@3/sd@0,0
       1. c2t1d0 <T300 cyl 34145 alt 2 hd 24 sec 128>
          /pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0
          ssd@w50020f2300000225,0
       2. c2t1d1 <T300 cyl 34145 alt 2 hd 20 sec 128>
          /pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0/
          ssd@w50020f2300000225,1
       3. c2t1d2 <SUN-T300-0117 cyl 34145 alt 2 hd 32 sec 128>
          /pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0/
          ssd@w50020f2300000225,2
       4. c2t1d3 <SUN-T300-0117 cyl 34145 alt 2 hd 32 sec 128>
          /pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0/
          ssd@w50020f2300000225,3
       5. c3t2d0 <T300 cyl 34145 alt 2 hd 24 sec 128>
          /pci@1f,2000/SUNW,qlc@1/fp@0,0/ssd@w50020f23000001f6,0
       6. c3t2d1 <T300 cyl 34145 alt 2 hd 20 sec 128>
          /pci@1f,2000/SUNW,qlc@1/fp@0,0/ssd@w50020f23000001f6,1
       7. c3t2d2 <SUN-T300-0117 cyl 34145 alt 2 hd 32 sec 128>
          /pci@1f,2000/SUNW,qlc@1/fp@0,0/ssd@w50020f23000001f6,2
       8. c3t2d3 <SUN-T300-0117 cyl 34145 alt 2 hd 32 sec 128>
          /pci@1f,2000/SUNW,qlc@1/fp@0,0/ssd@w50020f23000001f6,3
Specify disk (enter its number):
```

```
% metadb
          flags           first blk       block count
     a m  pc luo       16           1034          /dev/dsk/c2t1d0s2
     a    pc luo       1050         1034          /dev/dsk/c2t1d0s2
     a    pc luo       16           1034          /dev/dsk/c2t1d1s2
     a    pc luo       1050         1034          /dev/dsk/c2t1d1s2
```

```
% metadb
     flags           first blk       block count
     a m  pc luo       16           1034          /dev/dsk/c2t1d0s2
     a    pc luo       1050         1034          /dev/dsk/c2t1d0s2
     a    pc luo       16           1034          /dev/dsk/c2t1d1s2
     a    pc luo       1050         1034          /dev/dsk/c2t1d1s2
```

```
% metastat -p
d14 -m d10 d11 1
d10 1 2 c2t1d2s1 c2t1d2s6 -i 32b
d11 1 2 c2t1d3s1 c2t1d3s6 -i 32b
```

```
% metastat
d14: Mirror
    Submirror 0: d10
      State: Okay
    Submirror 1: d11
      State: Okay
    Pass: 1
    Read option: roundrobin (default)
    Write option: parallel (default)
    Size: 524288 blocks
d10: Submirror of d14
    State: Okay
    Size: 524288 blocks
    Stripe 0: (interlace: 32 blocks)
       Device            Start Block  Dbase State       Hot Spare
        c2t1d2s1                   0     No    Okay
        c2t1d2s6                   0     No    Okay
d11: Submirror of d14
    State: Okay
    Size: 524288 blocks
    Stripe 0: (interlace: 32 blocks)
       Device            Start Block  Dbase State       Hot Spare
        c2t1d3s1                   0     No    Okay
        c2t1d3s6                   0     No    Okay
```

5. **Unconfigure SDS without losing the data.**

   a. **Clear the submirrors and mirror devices.**

```
% metaclear d14
d14: Mirror is cleared
% metaclear d11
d11: Contact/Stripe is cleared
% metaclear d10
d10: Contact/Stripe is cleared
```

   a. **Clear the meta database replicas.**

```
% metadb -d -f -c 2 /dev/dsk/c2t1d0s7 /dev/dsk/c2t1d1s7
```

   Remove or comment out the `/etc/lvm/md.tab` entries.

```
#        d10             1 2 /dev/dsk/c2t1d2s1 /dev/dsk/c2t1d2s6
#        d11             1 2 /dev/dsk/c2t1d3s1 /dev/dsk/c2t1d3s6
#        d14             -m d10 d11
```

6. **Enable the Sun StorEdge Traffic Manager software as described in *Post Installation Procedure* in Chapter "Post Installation Tasks."**

7. **Determine the Sun StorEdge Traffic Manager software device name to pre-STMS device name mapping.**

The output of `format`, `luxadm probe`, and `display` gives the Sun StorEdge Traffic Manager software paths for each device. In this case the output looks:

```
% format
AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
          /pci@1f,4000/scsi@3/sd@0,0
       1. c4t60020F20000002253B220F99000D348Cd0 <SUN-T300-0117
          cyl 34145 alt 2 hd 32 sec 128>
          /scsi_vhci/ssd@g60020f20000002253b220f99000d348c
       2. c4t60020F20000002253B220FC000086944d0 <SUN-T300-0117
          cyl 34145 alt 2 hd 32 sec 128>
          /scsi_vhci/ssd@g60020f20000002253b220fc000086944
       3. c4t60020F20000002253B220FD400071CD8d0 <SUN-T300-0117
          cyl 34145 alt 2 hd 32 sec 128>
          /scsi_vhci/ssd@g60020f20000002253b220fd400071cd8
       4. c4t60020F20000002253B22101D00029FC9d0 <SUN-T300-0117
          cyl 34145 alt 2 hd 32 sec 128>
          /scsi_vhci/ssd@g60020f20000002253b22101d00029fc9
Specify disk (enter its number): ^D
```

```
% luxadm probe
No Network Array enclosures found in /dev/es

Found Fibre Channel device(s):
  Node WWN:50020f20000001f6  Device Type:Disk device
   Logical Path:/dev/rdsk/c4t60020F20000002253B22101D00029FC9d0s2
  Node WWN:50020f2000000225  Device Type:Disk device
   Logical Path:/dev/rdsk/c4t60020F20000002253B220FD400071CD8d0s2
  Node WWN:50020f20000001f6  Device Type:Disk device
   Logical Path:/dev/rdsk/c4t60020F20000002253B220FC000086944d0s2
  Node WWN:50020f2000000225  Device Type:Disk device
   Logical Path:/dev/rdsk/c4t60020F20000002253B220F99000D348Cd0s2

For each entry in luxadm probe, get the device path information.
```

```
% luxadm display
/dev/rdsk/c4t60020F20000002253B22101D00029FC9d0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c4t60020F20000002253B22101D00029FC9d0s2
  Status(Port A):       O.K.
  Status(Port B):       O.K.
  Vendor:               SUN
  Product ID:           T300
  WWN(Node):            50020f20000001f6
  WWN(Port A):          50020f23000001f6
  WWN(Port B):          50020f2300000225
  Revision:             0117
  Serial Num:           054921
  Unformatted capacity: 68294.000 MBytes
  Write Cache:          Enabled
  Read Cache:           Enabled
    Minimum prefetch:   0x0
    Maximum prefetch:   0x0
  Device Type:          Disk device
  Path(s):
  /dev/rdsk/c4t60020F20000002253B22101D00029FC9d0s2
  /devices/scsi_vhci/ssd@g60020f20000002253b22101d00029fc9:c,raw
   Controller         /devices/pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0
     Device Address     50020f2300000225,3
     Class              secondary
     State              STANDBY
    Controller          /devices/pci@1f,2000/SUNW,qlc@1/fp@0,0
     Device Address     50020f23000001f6,3
     Class              primary
     State              ONLINE
```

```
% luxadm display
/dev/rdsk/c4t60020F20000002253B220FD400071CD8d0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c4t60020F20000002253B220FD400071CD8d0s2
  Status(Port A):       O.K.
  Status(Port B):       O.K.
  Vendor:               SUN
  Product ID:           T300
  WWN(Node):            50020f2000000225
  WWN(Port A):          50020f23000001f6
  WWN(Port B):          50020f2300000225
  Revision:             0117
  Serial Num:           054911
  Unformatted capacity: 68294.000 MBytes
  Write Cache:          Enabled
 Read Cache:           Enabled
    Minimum prefetch:   0x0
    Maximum prefetch:   0x0
  Device Type:          Disk device
  Path(s):
  /dev/rdsk/c4t60020F20000002253B220FD400071CD8d0s2
  /devices/scsi_vhci/ssd@g60020f20000002253b220fd400071cd8:c,raw
   Controller         /devices/pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0
     Device Address     50020f2300000225,2
     Class              secondary
     State              ONLINE
    Controller          /devices/pci@1f,2000/SUNW,qlc@1/fp@0,0
     Device Address     50020f23000001f6,2
     Class              primary
     State              STANDBY
```

```
% luxadm display
/dev/rdsk/c4t60020F20000002253B220F99000D348Cd0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c4t60020F20000002253B220F99000D348Cd0s2
  Status(Port A):       O.K.
  Status(Port B):       O.K.
  Vendor:               SUN
  Product ID:           T300
  WWN(Node):            50020f2000000225
  WWN(Port A):          50020f2300000225
  WWN(Port B):          50020f23000001f6
  Revision:             0117
  Serial Num:           054910
  Unformatted capacity: 68294.000 MBytes
  Write Cache:          Enabled
  Read Cache:           Enabled
    Minimum prefetch:   0x0
    Maximum prefetch:   0x0
  Device Type:          Disk device
  Path(s):
  /dev/rdsk/c4t60020F20000002253B220F99000D348Cd0s2
  /devices/scsi_vhci/ssd@g60020f20000002253b220f99000d348c:c,raw
   Controller         /devices/pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0
     Device Address     50020f2300000225,0
     Class              primary
     State              ONLINE
    Controller          /devices/pci@1f,2000/SUNW,qlc@1/fp@0,0
     Device Address     50020f23000001f6,0
     Class              secondary
     State              STANDBY
```

Now, correlate the non-STMS device information collected in Step 1 by matching the `luxadm display` paths output. Identify the appropriate lun.

For example, to find c2t1d0 device under the Sun StorEdge Traffic Manager software.

c2t1d0 is lun 0 of the T3 partner pair as seen from the format in Step 1.

```
1. c2t1d0 <T300 cyl 34145 alt 2 hd 24 sec 128>
/pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0/ssd@w50020f2300000225,0
```

From `luxadm display` of each device above, looking at the Paths, controller and device address fields.

Path(s):

```
/dev/rdsk/c4t60020F20000002253B220F99000D348Cd0s2
/devices/scsi_vhci/ssd@g60020f20000002253b220f99000d348c:c,raw
Controller /devices/pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0
Device Address 50020f2300000225,0
```

Thus, c2t1d0 is now /dev/rdsk/c4t60020F20000002253B220F99000D348Cd0
c2t1d1 LUN 1 is now /dev/rdsk/c4t60020F20000002253B220FC000086944d0
c2t1d2 LUN 2 is now /dev/rdskc4t60020F20000002253B220FD400071CD8d0
And, c2t1d3 LUN 3 is /dev/rdsk/c4t60020F20000002253B22101D00029FC9d0

8. Reconfigure SDS under the Sun StorEdge Traffic Manager software.

a. Recreate metadb database replicas with new device names.

```
% metadb
flags           first blk       block count
a m  pc luo       16              1034           /dev/dsk/c2t1d0s2
a    pc luo       1050            1034           /dev/dsk/c2t1d0s2
a    pc luo       16              1034           /dev/dsk/c2t1d1s2
a    pc luo       1050            1034           /dev/dsk/c2t1d1s2
```

```
% metadb -a -f -c 2
/dev/rdsk/c4t60020F20000002253B220F99000D348Cd0s2
/dev/rdsk/c4t60020F20000002253B220FC000086944d0s2
```

b. Similarly, identify the metadevices with new device names.

```
% metastat -p
d14 -m d10 d11 1
d10 1 2 c2t1d2s1 c2t1d2s6 -i 32b
d11 1 2 c2t1d3s1 c2t1d3s6 -i 32b
```

c. Recreate submirrors and mirror.

```
% metainit d10 1 2 /dev/dsk/
c4t60020F20000002253B220FD400071C8d0s1
/dev/dsk/c4t60020F20000002253B220FD400071CD8d0s6
d10: Concat/Stripe is setup
```

```
% metainit d11 1 2 /dev/dsk/
c4t60020F20000002253B22101D00029F9d0s1
/dev/dsk/c4t60020F20000002253B22101D00029FC9d0s6
d11: Concat/Stripe is setup
```

```
% metanit d14
metainit: d14: WARNING: This form of metainit is not recommended.
The submirrors may not have the same data.
Please see ERRORS in metainit(1M) for additional information.
d14: Mirror is setup
```

Verify that all the data is now available in the mirror device. d14 is the same as the data prior to enabling the Sun StorEdge Traffic Manager software on the same mirror.

CHAPTER 4

Deployment Scenarios

This chapter shows two different types of configurations: First, a system that uses STMS, and second a system that does not use STMS.

Topics in this chapter include:

- "About Traffic Manager Configuration Samples" on page 43
- "Configuration without Traffic Manager" on page 44
- "Configuration with Traffic Manager" on page 45

About Traffic Manager Configuration Samples

STMS device enumeration is slightly different from legacy devices in that only one device path is shown per device regardless of the number of paths.

Storage devices such as A3x00FC have their own multipathing solution. STMS can co-exist with such storage devices. However, these devices will not be enumerated under STMS and will work the same as if STMS were not installed.

Here we look at A5x00 and T3 along with the utilities.

Case Study: A host has the following storage attached.

- A5x00 with both A and B loop connected.
- T3 partner pair with 4 LUNs.

See FIGURE 1-3 and FIGURE 1-4 in Chapter 1 for illustrations of the two configurations described in this chapter.

Configuration without Traffic Manager

Before STMS was installed and configured, T3 mp_support was set to rw format. The output is now different from the configuration that included STMS.

Utility: format

The output for a configuration that does not use STMS is as follows:

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
       0.  c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
           /pci@1f,4000/scsi@3/sd@0,
       1. c2t3d0 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
           /pci@1f,4000/SUNW,qlc@2/fp@0,0/ssd@w50020f23000042d4,0
       2. c2t3d1 <SUN-T300-0116 cyl 34145 alt 2 hd 24 sec 128>
           /pci@1f,4000/SUNW,qlc@2/fp@0,0/ssd@w50020f23000042d4,1
       3. c2t3d2 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
           /pci@1f,4000/SUNW,qlc@2/fp@0,0/ssd@w50020f23000042d4,2
       4. c2t3d3 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
           /pci@1f,4000/SUNW,qlc@2/fp@0,0/ssd@w50020f23000042d4,3
       5. c3t4d0 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
           /pci@1f,4000/SUNW,qlc@4/fp@0,0/ssd@w50020f2300003fad,0
       6. c3t4d1 <SUN-T300-0116 cyl 34145 alt 2 hd 24 sec 128>
           /pci@1f,4000/SUNW,qlc@4/fp@0,0/ssd@w50020f2300003fad,1
       7. c3t4d2 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
           /pci@1f,4000/SUNW,qlc@4/fp@0,0/ssd@w50020f2300003fad,2
       8. c3t4d3 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
           /pci@1f,4000/SUNW,qlc@4/fp@0,0/ssd@w50020f2300003fad,3
       9. c4t68d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
           /pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0
           ssd@w22000020371a1862,0
       10. c5t68d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
           /pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,
           ssd@w21000020371a1862,0
Specify disk (enter its number): ^D
```

Configuration with Traffic Manager

After STMS is installed and configured, T3 mp_support is set to “mpxio” (see “Chapter Installing STMS.”

Utility: format

The output of `format` looks as follows:

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
       0.  c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
           /pci@1f,4000/scsi@3/sd@0,0
       1.  c6t60020F20000042D43ADCBC4E000C41E2d0 <SUN-T300-0117
           cyl 34145 alt 2 hd 24 sec 128>
           /scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2
       2.  c6t60020F20000042D43B0E926A000AA3FCd0 <SUN-T300-0116
           cyl 34145 alt 2 hd 24 sec 128>
           /scsi_vhci/ssd@g60020f20000042d43b0e926a000aa3fc
       3.  c6t60020F20000042D43B2753510008C9DFd0 <SUN-T300-0117
           cyl 34145 alt 2 h d 24 sec 128>
           /scsi_vhci/ssd@g60020f20000042d43b2753510008c9df
       4.  c6t60020F20000042D43B275377000877DDd0 <SUN-T300-0117
           cyl 34145 alt 2 h d 24 sec 128>
           /scsi_vhci/ssd@g60020f20000042d43b275377000877dd
       5.  c6t20000020371A1862d0 <SUN36G cyl 24620 alt 2 hd 27 sec
107>
           /scsi_vhci/ssd@g20000020371a1862
Specify disk (enter its number): ^D
```

Consider the following notes:

1. Devices enumerated under STMS will have /scsi_vhci/ssd entry.
2. The first four scsi_vhci entries are the four T3 LUNs now have a very long name which is Global Unique Identifier (GUID). Also, only one entry per lun is visible instead of two paths.
3. Next scsi_vhci entries A5x00 devices have its WWN (World Wide Number) and only one device entry is created for both of A and B loops.

4. There is only one controller number assigned for all the devices encapsulated under STMS.

5. Use `luxadm` display to identify the mapping between device entries without STMS and the device entries with STMS.

Utility: luxadm probe

The `luxadm probe` now shows the WWN for A5x00 and GUID for T3.

```
# luxadm probe
Found Enclosure:
SENA Name:f  Node WWN:50800200000777d0  Logical Path:/dev/es/ses0
Found Fibre Channel device(s):
  Node WWN:50020f20000042d4  Device Type:Disk device
   Logical Path:/dev/rdsk/c6t60020F20000042D43B275377000877DDd0s2
  Node WWN:50020f20000042d4  Device Type:Disk device
   Logical Path:/dev/rdsk/c6t60020F20000042D43B2753510008C9DFd0s2
  Node WWN:50020f20000042d4  Device Type:Disk device
   Logical Path:/dev/rdsk/c6t60020F20000042D43B0E926A000AA3FCd0s2
  Node WWN:50020f20000042d4  Device Type:Disk device
   Logical Path:/dev/rdsk/c6t60020F20000042D43ADCBC4E000C41E2d0s2
```

Utility: luxadm display

The `luxadm display` has been enhanced for STMS. For each entry in the `format` or `luxadm probe`, `luxadm` display output indicates the following:

1. The number of paths to the storage device.

2. The mapping of the paths prior to STMS and after STMS (under Path(s) controller and device address).

3. The state and type of each path.

   a. State of each path.

      i. ONLINE indicates the active path(s) through which IO is going to the device. For more than one ONLINE path, STMS will use load balancing (like round robin scheme) or IO to the device.

ii. STANDBY indicates the path is available if an ONLINE path fails or is switched to another state. There can be many STANDBY paths. If a STANDBY path is chosen to be the active path for routing IO, its status will be changed to ONLINE.

iii. OFFLINE indicates the path(s) was previously existing but is not available now.

b. Class type for each path.

i. PRIMARY: This path is the preferred path for routing IO.

ii. SECONDARY: This path is the next priority path after PRIMARY.

If a Sun StorEdge T3 partner pair configuration is used, two paths for each lun exist. One path is ONLINE and the other is STANDBY. In this configuration, IO to the lun is active on its ONLINE path. If this path fails, the STANDBY path becomes the ONLINE path. However, if the first path later becomes available, it will go into STANDBY mode instead of ONLINE, thus saving an expensive failover operation. Use `luxadm` failover command to bring the restored path to ONLINE (see `luxadm` man pages).

As can be seen from the following two `luxadm display` output of some T3 LUNs, PRIMARY can be STANDBY and SECONDARY can be ONLINE for one lun, whereas SECONDARY is STANDBY and PRIMARY is ONLINE for another LUN.

```
# luxadm display
/dev/rdsk/c6t60020F20000042D43B275377000877DDd0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c6t60020F20000042D43B275377000877DDd0s2
  Status(Port A):       O.K.
  Status(Port B):       O.K.
  Vendor:               SUN
  Product ID:           T300
  WWN(Node):            50020f20000042d4
  WWN(Port A):          50020f2300003fad
  WWN(Port B):          50020f23000042d4
  Revision:             0117
  Serial Num:           Unsupported
  Unformatted capacity: 51220.500 MBytes
  Write Cache:          Enabled
  Read Cache:           Enabled
    Minimum prefetch:   0x0
    Maximum prefetch:   0x0
  Device Type:          Disk device
  Path(s):
  /dev/rdsk/c6t60020F20000042D43B275377000877DDd0s2
  /devices/scsi_vhci/ssd@g60020f20000042d43b275377000877dd:c,raw
   Controller           /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
     Device Address     50020f23000042d4,3
           Class              secondary
     State              ONLINE
   Controller           /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
     Device Address     50020f2300003fad,3
     Class              primary
     State              STANDBY
```

```
# luxadm display
/dev/rdsk/c6t60020F20000042D43B0E926A000AA3FCd0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c6t60020F20000042D43B0E926A000AA3FCd0s2
  Status(Port A):       O.K.
  Status(Port B):       O.K.
  Vendor:               SUN
  Product ID:           T300
  WWN(Node):            50020f20000042d4
  WWN(Port A):          50020f23000042d4
  WWN(Port B):          50020f2300003fad
  Revision:             0117
  Serial Num:           Unsupported
  Unformatted capacity: 51220.500 MBytes
  Write Cache:          Enabled
  Read Cache:           Enabled
    Minimum prefetch:   0x0
    Maximum prefetch:   0x0
  Device Type:          Disk device
  Path(s):
  /dev/rdsk/c6t60020F20000042D43B0E926A000AA3FCd0s2
  /devices/scsi_vhci/ssd@g60020f20000042d43b0e926a000aa3fc:c,raw
Controller           /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
     Device Address     50020f23000042d4,1
     Class              primary
     State              ONLINE
Controller           /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
     Device Address     50020f2300003fad,1
     Class              secondary
     State              STANDBY
```

In the case of A5x00 in the current configuration, there are two paths. Both are ONLINE and IO is load balanced on both the paths. If a path fails, the second path continues the IO. If the failed path comes back, it will be in ONLINE state and starts participating in IO transfer.

```
# luxadm display /dev/rdsk/c6t20000020371A1862d0s2
DEVICE PROPERTIES for disk: /dev/rdsk/c6t20000020371A1862d0s2
  Status(Port A):       O.K.
  Status(Port B):       O.K.
  Vendor:               SEAGATE
  Product ID:           ST136403FSUN36G
  WWN(Node):            20000020371a1862
  WWN(Port A):          21000020371a1862
  WWN(Port B):          22000020371a1862
  Revision:             114A
  Serial Num:           LT0187150000
  Unformatted capacity: 34732.891 MBytes
  Read Cache:           Enabled
    Minimum prefetch:   0x0
    Maximum prefetch:   0xffff
  Location:             In the enclosure named: f
  Device Type:          Disk device
  Path(s):
  /dev/rdsk/c6t20000020371A1862d0s2
  /devices/scsi_vhci/ssd@g20000020371a1862:c,raw
   Controller         /devices/pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0
      Device Address     22000020371a1862,0
      Class              primary
            State              ONLINE
   Controller         /devices/pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,0
      Device Address     21000020371a1862,0
      Class              primary
      State              ONLINE
```

Utility: luxadm failover

The `luxadm failover` command is used to failover LUNs from primary to secondary paths and vice versa.

In case of Sun StorEdge T3 partner pair, when a failover happens, for example, due to cable pull to one T3 controller, the LUNs owned by that controller are failed over to the alternate controller.

However, in the current implementation of STMS when the fault is rectified, the LUNs do not failover to the original configuration automatically.

A `luxadm failover` sub-command has to be issued to perform failover to the original configuration.

Example of Failover:

1. Original state of a Sun StorEdge T3 LUN is obtained by `luxadm display` command.

```
# luxadm display /dev/rdsk/c5t60020F20000042D43ADCBC4E000C41Ed0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
  Status(Port A):       O.K.
  Status(Port B):       O.K.
  Vendor:               SUN
  Product ID:           T300
  WWN(Node):            50020f20000042d4
  WWN(Port A):          50020f23000042d4
  WWN(Port B):          50020f2300003fad
  Revision:             0117
  Serial Num:           Unsupported
  Unformatted capacity: 51220.500 MBytes
  Write Cache:          Enabled
  Read Cache:           Enabled
    Minimum prefetch:   0x0
    Maximum prefetch:   0x0
  Device Type:          Disk device
  Path(s):
  /dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
  /devices/scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2:c,raw
   Controller           /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
     Device Address     50020f23000042d4,0
     Class              primary
     State              ONLINE
   Controller           /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
     Device Address     50020f2300003fad,0
     Class              secondary
     State              STANDBY
```

Note – The primary path is ONLINE and the secondary path is STANDBY.

2. Cable is pulled from the T3 controller 50020f23000042d4. Failover is triggered and the primary path is OFFLINE and secondary path is ONLINE. LUN status now is degraded.

```
# luxadm display /dev/rdsk/c5t60020F20000042D43ADCBC4E000C41Ed0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
  ..
  Path(s):
  /dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
  /devices/scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2:c,raw
   Controller           /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
     Device Address     50020f23000042d4,0
     Class              primary
     State              OFFLINE
   Controller           /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
     Device Address     50020f2300003fad,0
     Class              secondary
     State              ONLINE
```

3. Cable is reinserted to T3 controller 50020f23000042d4. The device state become optimal but, the failover is not triggered. The Primary path comes up as STANDBY and still the secondary path is ONLINE.

```
# luxadm display /dev/rdsk/c5t60020F20000042D43ADCBC4E000C41Ed0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
..
  Path(s):
  /dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
  /devices/scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2:c,raw
   Controller           /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
     Device Address     50020f23000042d4,0
     Class              primary
     State              STANDBY
   Controller           /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
     Device Address     50020f2300003fad,0
     Class              secondary
     State              ONLINE
```

4. Enter the `luxadm failover` command to failover to primary.

```
# luxadm failover primary
/dev/rdsk/c5t60020F20000042D43ADCBCE000C41E2d0s2
#

Now this triggers the failover and primary path becomes ONLINE and
secondary path becomes STANDBY which is equivalent to the original
state in Step 1.

# luxadm display /dev/rdsk/c5t60020F20000042D43ADCBC4E000C41Ed0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
 ..
  Path(s):
  /dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
  /devices/scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2:c,raw
   Controller           /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
     Device Address     50020f23000042d4,0
     Class              primary
     State              ONLINE
   Controller           /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
     Device Address     50020f2300003fad,0
     Class              secondary
     State              STANDBY
```

For more details on the `luxadm failover` command, please refer to the man pages.

CHAPTER 5

Troubleshooting STMS Software

This chapter provides solutions to potential problems that may occur while running Sun StorEdge Traffic Manager software.

Topics in this chapter include:

- "STMS Not Running Properly" on page 56
- "STMS Running Properly, but luxadm display and luxadm failover are Failing" on page 56
- "Sun StorEdge T3 Arrays Do Not Show" on page 57
- "System Failed During Boot with scsi_vhci Attachment" on page 57
- "Group of A5x00s Connected, But They Still Appear Under Physical Paths in Format" on page 58
- "System and Error Messages" on page 58

STMS Not Running Properly

The first item to check is if STMS has been installed correctly.

Verify STMS drivers that is, scsi_vhci and STMS, are loaded with the help of the `modinfo` command. (See the `modinfo` man pages.)

```
# modinfo | grep mpxio
23 102193e5   84cb   -   1  mpxio (MDI Library)
# modinfo | grep scsi_vhci
121 781ea000   6a20 225   1 scsi_vhci (Sun Multiplexed SCSI vHCI)
```

If `modinfo` output does not show the drivers are loaded, verify that the following binaries have been loaded in the proper directories.

- `/kernel/drv/scsi_vhci, /kernel/drv/sparcv9/scsi_vhci` (64 bit)
- `/kernel/misc/mpxio, /kernel/misc/sparcv9/mpxio` (64-bit)
- `/kernel/drv/scsi_vhci.conf`

If these binaries are not present in the specified directories, then STMS did not install properly. Please repeat the installation process. Make sure you are logged in as superuser when installing the software.

STMS Running Properly, but `luxadm display` and `luxadm failover` are Failing

Ensure that all your qlc pHCIs are seen using the `luxadm qlgc` command. If you don't see one or more of the pHCIs, then their Fcode is out of date. An alternate way to confirm this is by examining the device path. Notice the "scsi@4" in the example below. Make sure that none of the device paths looks like the following:

`/devices/pci@9d,600000/pci@1/scsi@4/fp@0,0:devctl`

To correct this problem, download the latest Fcode software.

Sun StorEdge T3 Arrays Do Not Show

A Sun StorEdge T3 array can be correctly configured for STMS by using the latest firmware revision. The latest at the time this document was written is 1.17a. Be sure to set the master settings as follows:

```
t31> sys mp_support mpxio
```

See the *Sun StorEdge Traffic Manager Release Notes: For the Solaris Operating Environment* for more details about the system requirements.

While checking the Sun StorEdge T3 array setup, confirm that the LPC version is also current. Once the device is configured, it is advisable to perform a reconfigure boot or equivalent.

System Failed During Boot with `scsi_vhci` Attachment

It is most likely due to an incomplete installation. The installation failed to provide an entry for the scsi_vhci in the `name_to_major` database. If you do not see any installation log errors, call the support.

Group of A5x00s Connected, But They Still Appear Under Physical Paths in Format

Consider the following:

- Check whether the A5x00 class devices are not connected to a STMS supported HBA mentioned in Chapter “Requirements and Environments.”.
- Check if STMS is disabled under `/kernel/drv/scsi_vhci.conf`.
- Check if the system is booting from the A5x00 disks (booting from StorEdge T3 is not supported). When you boot from a STMS device, all devices under the pHCI with the boot device get enumerated under scsi_vhci.
- Make sure the pHCI is disabled under STMS from `/kernel/drv/qlc.conf`.

System and Error Messages

The following messages might appear in the course of operation.

- When the automatic failback feature is enabled through the configuration file, you should receive the following message:

```
Auto-failback capability enabled through scsi_vhci.conf file.
```

- If automatic failback succeeds, the following message is logged:

```
Auto failback operation succeeded for devices.
```

- If automatic failback fails, the following message is logged:

```
Auto failback operation failed for device.
```

- Externally initiated failover of a Sun StorEdge T3 or T3+ array has been observed.

```
Waiting for externally initiated failover to complete
```

```
path target address %s is now ONLINE because of an externally
initiated failover
```

```
path target address %s is now STANDBY because of an externally
initiated failover %s (GUID %s)
```

- Failover has been prevented by the cause indicated

```
Failover operation failed for device: due to an internal error
```

```
Failover operation failed for device: Invalid path-class
```

```
Failover operation failed for device: no alternate paths found
```

```
Failover operation completed successfully for device: failed over
from _ to _
```

```
T3 failover failed, couldn’t transport packet
```

```
T3 failover failed: timed out
```

```
T3 failover failed: timed out waiting for path to become active
```

```
T3 failover failed; sense key:%x, ASC: %x, ASCQ:%x
```

Glossary

A

Array A Sun StorEdge T3 array that contains an internal RAID controller and nine disk drives with Fibre Channel connectivity to the data host.

C

Client Device A multipath device attached beneath an instance of a vHCI driver.

F

Failover Routing I/O requests through alternate paths when a failure is detected on the primary or active HCI.

G

Global Unique Identifier (GUID) While the definition is command set specific, a GUID generally is an identifier that has been defined unique, as in no 2 devices will have the same ID.

H

HCI The host controller interface (HCI) consists of the software and hardware required to connect the storage array to the server.

Hot Standby A pHCI which is configured to not transport data traffic unless all primary interfaces have failed. A pHCI configured in this way may carry a small amount of traffic for latent fault detection to ensure that the standby is operational.

L

Load Balancing The ability to route I/O requests through different paths for better utilization of host controller resources.

Logical Unit Number (LUN) One or more drives that are grouped into a unit; they are also called a "volume."

M

MDI The multipath driver interface (MDI) is a new set of interfaces for HCI drivers.

P

Packet Transport The SCSI vHCI driver interfaces with the Traffic Manager architecture to select a suitable path for routing a request.

Partner Group A pair of interconnected arrays.

Path Another term for a pHCI transport service which is being managed by a vHCI multipathing nexus driver. The vHCI driver selects a path from a list of available paths.

Path Set Set of paths which share a common set of attributes. Path sets are created automatically as part of the device discovery and enumeration process. Path sets may also be instantiated and managed by applications.

Physical Host Controller Interface (pHCI) (a.k.a. HCI) An instance of bus nexus device providing command set specific transport and naming services for devices attached to it via an I/O interconnect.

R

Repair Detection Detecting that a path has begun to function correctly after having previously failed.

S

SCSA Sun compliant SCSI adapter.

Switchover Routing I/O through other valid paths through an administrative utility. Similar to the Alternate Pathing disk driver developed by DHPG.

V

vHCI A “virtual” pseudo nexus driver under which multipath client devices are attached. The vHCI driver is responsible for routing transport requests from the client devices to the appropriate pHCI devices which are providing transport services to the device.

Volume One or more drives that are grouped together into a unit; they are also called “LUN.”

W

World Wide Name (WWN) Unique number assigned to each device on a Fibre Channel loop.