



Sun StorEdge™ SAN Foundation Software 4.4 Configuration Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 817-3672-11
March 2004

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

Copyright 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun StorEdge, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Netscape Navigator is a trademark or registered trademark of Netscape Communications Corporation in the United States and other countries.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuelle relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Sun StorEdge, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

Netscape Navigator est une marque de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Adobe PostScript

Contents

Preface	xi
1. Features and Guidelines	1
Fabric Bootability	2
SNIA (FC-HBA) and FCSM Packages	2
Zones and Ports	6
Zone Types	6
Port Types	7
Guidelines	8
Cascading Guidelines	8
IPFC Guidelines	9
Multipathing Guidelines	10
2. Configuring the Switch Ethernet Port	11
Using RARP to Change the IP Address	12
▼ To Prepare the File System for RARP	12
▼ To Set the Network IP Address With RARP	14
3. Booting From Fibre Channel Storage Devices	15
Special Bootability Considerations	16
Determining Which Boot Method to Use	17

Net Install Patching Method 18

Overview Of the Net Install Method 19

- ▼ To Set Up the Boot or Install Server 20
- ▼ To Modify the Boot Image 21
- ▼ To Set Up the Client and Reboot 22

Dump and Restore Method 24

Overview Of the Dump and Restore Method 24

Partitioning the New Boot Disk the Same As the Temporary Boot Disk 25

- ▼ To Prepare To Partition the New Disk 25
- ▼ To Record the Partition Layout 26
- ▼ To Change to the New Boot Disk 29
- ▼ To Specify Slices on the New Boot Disk 31
- ▼ To Label the New Boot Disk 33
- ▼ To Create File Systems on the New Boot Disk 34

Creating the New Boot Files 35

- ▼ To Copy the Boot Block and Root File System Contents to the New Boot Disk 35
- ▼ To Update the `vfstab` File 37
- ▼ To Copy the Contents of Non-Root File Systems Onto the New Boot Disk 38
- ▼ To Specify the New Boot Disk as the Boot Device 39

4. Adding and Removing SAN Devices 43

Special Considerations 44

Port Choices 44

Multipathing 45

Adding and Removing Devices 45

Adding a Storage Device 45

- ▼ To Add a Storage Device 46

Removing a Storage Device 48

- ▼ To Remove a Storage Device 49

5. Fabric Device Node Recognition 51

Creating and Removing Individual Device Nodes Without Multipathing Enabled 52

- ▼ To Detect Fabric Devices Visible on a Host 53

Ensuring LUN Level Information Is Visible 54

- ▼ To Configure a Fabric Device Without Multipathing Enabled On the Host 55

- ▼ To Unconfigure a Fabric Device 59

Creating and Removing Multiple Device Nodes Without Multipathing Enabled 61

- ▼ To Configure All Fabric Devices on a Fabric-Connected Host Port 61

- ▼ To Unconfigure All Fabric Devices on a Fabric-Connected Host Port 63

Creating and Removing Individual Device Nodes With Multipathing Enabled 65

- ▼ To Detect Fabric Devices Visible to a Host 66

Ensuring LUN Level Information Is Visible 67

- ▼ To Configure Fabric-Connected Device Nodes With Multipathing Enabled Devices 68

- ▼ To Unconfigure a Fabric Device Associated With Multipathing Enabled Devices 72

- ▼ To Unconfigure One Path to a Multipathed Device 74

Creating and Removing Multiple Device Nodes With Multipathing Enabled 77

- ▼ To Configure All Fabric-Connected Devices With Multipathing Enabled 78

- ▼ To Unconfigure All Fabric-Connected Devices With Multipathing Enabled 83

6. IPFC Management 85

Loading IPFC 86

- Determining Fibre Channel Adapter Port Instances 86

- ▼ To Plumb an IPFC Instance 89

IPFC Invocation, Configuration and Usage 91

- ▼ To Start a Network Interface Manually 91
- ▼ To Configure the Host for Automatic Plumbing Upon Reboot 92

IPFC Usage 92

7. Dynamic Reconfiguration On the SAN 93

DR and Fabric Devices 93

- ▼ To Remove a Fabric Device Before DR 94
- ▼ To Maintain a Fabric Device Configuration With DR 94
- ▼ To Reconfigure Fabric Devices With DR 94
- ▼ To Reconfigure the Sun Enterprise 10000 Server With a Fabric Connection 94

Glossary 97

Index 101

Figures

FIGURE 4-1 Decision Tree For Adding Fabric Devices 47

FIGURE 4-2 Decision Tree For Removing Fabric Devices 48

Tables

TABLE 1-1	Supported and Unsupported SNIA Interfaces	3
TABLE 1-2	Sun Switch Port Types	7
TABLE 1-3	IPFC (NFS/NAS and SNDR)	9
TABLE 4-1	Port Types and For Storage Devices	44
TABLE 6-1	PCI Single Fibre Channel Network Adapter <code>/etc/path_to_inst</code> Device Path Entry	89

Preface

The *Sun StorEdge SAN Foundation Software 4.4 Configuration Guide* describes how to configure various devices in a Sun StorEdge SAN Foundation software environment.

This guide is intended for system, storage, and network administrators who create and maintain storage area networks.

Before You Read This Book

Before you read this book, review the late-breaking information described in the *Sun StorEdge SAN Foundation Software 4.4 Release Notes* and complete the software installation for your SAN with the help of the *Sun StorEdge SAN Foundation Software 4.4 Installation Guide*.

How This Book is Organized

- Chapter 1 provides an overview of features and guidelines.
- Chapter 2 provides instructions on configuring ports on selected switches.
- Chapter 3 explains how to configure booting from Fibre Channel storage devices.
- Chapter 4 explains how to add or remove devices from your SAN.
- Chapter 5 instructs how to configure fabric device node recognition with or without multipathing enabled.
- Chapter 6 explains how to load IP over Fibre Channel (IPFC) drivers and how to invoke IPFC.
- Chapter 7 explains how to configure your SAN for dynamic reconfiguration.
- A glossary and index define terms and list where they can be found in this document.

Using UNIX Commands

This document may 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

TABLE P-1

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

TABLE P-2

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

TABLE P-3 Sun StorEdge SAN Foundation Release Related Documentation

Product	Application	Title	Part Number
Sun StorEdge SAN 4.4 Foundation Release	Documentation information	<i>Sun StorEdge SAN Foundation Software 4.4 Guide to Documentation</i>	817-3670-11

Contacting Sun Technical Support

If you need help installing or using this product, contact your service provider. If you have a support contract with Sun, call 1-800-USA-4SUN, or go to:

<http://www.sun.com/service/contacting/index.html>

Accessing Sun Documentation

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

<http://www.sun.com/documentation>

A broad selection of Sun system documentation is located at:

<http://www.sun.com/products-n-solutions/hardware/docs>

A complete set of Solaris documentation and many other titles are located at:

<http://docs.sun.com>

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 SAN Foundation Software 4.4
4 Configuration Guide, part number 817-3672-11

Features and Guidelines

Before using this manual, you should have successfully completed the Sun StorEdge SAN Foundation (SFS) hardware and software installation. The *Sun StorEdge SAN Foundation Software 4.4 Release Notes* identifies hardware on the Sun price list that you can include in your network. It also lists features and functionality that Sun supports.

This section provides related guidelines for creating a SAN with the SFS and contains the following topics:

- “Fabric Bootability” on page 2
- “SNIA (FC-HBA) and FCSM Packages” on page 2
- “Zones and Ports” on page 6
- “Guidelines” on page 8

Fabric Bootability

The 1 Gbit and 2 Gbit host bus adapters (HBAs) that support the SFS can now boot from fabric devices as well as non-fabric devices. Fabric topologies with Fibre Channel switches provide higher speed, more connections, and port isolation. A Fibre Channel port can be configured in a loop or fabric topology.

As Sun obsoletes or removes loop support with its Fibre Channel switches, some current bootable configurations will not work in the future. With support for fabric bootability, however, you can still boot from devices on your SAN as you use the new switches. For fabric bootability instructions, see “Dump and Restore Method” on page 24 in this guide and also the information about fabric boot in the *Sun StorEdge SAN Foundation Software 4.4 Installation Guide*.

SNIA (FC-HBA) and FCSM Packages

Note – SNIA (for example, “SNIA common library”) is now also known as FC-HBA.

The SAN Foundation release includes the Storage Networking Industry Association (SNIA) Fibre Channel host bus adapter (HBA) library. Some of these interfaces are provided through the Fibre Channel Switch Management (FCSM) driver. The SNIA library uses the FCSM driver to send Fibre Channel Common Transport (CT) requests to Fibre Channel switches. This allows client applications to get detailed information about the SAN topology and switch configurations.

The SNIA HBA application programming interfaces (APIs) enable you to manage Fibre Channel HBAs in your SAN. You can use the interfaces to access detailed information about the HBAs on a host and switches and storage devices connected to the SAN to comply with industry standards. Specifically, SNIA and FCSM enable you to integrate the Sun StorEdge Enterprise Storage Manager Topology Reporter (ESM) and other third-party SAN management tools with the Sun StorEdge SAN Foundation software required for devices that Sun supports.

The SNIA HBA API is divided into two functional components. The first is a Common Library, to which applications are linked. The second is one or more Vendor Specific Libraries (VSLs), which the Common Library dynamically loads to manage individual vendor HBAs.

Man pages for the common SNIA API are included in the Sun StorEdge SAN Foundation software. The formal specification (FC-MI) is available at <http://www.t11.org>. TABLE 1-1 outlines which APIs are included in the vendor library for this release:

TABLE 1-1 Supported and Unsupported SNIA Interfaces

SNIA API	Sun StorEdge SAN Foundation Software Support
HBA_GetVersion	Yes
HBA_LoadLibrary	Yes
HBA_FreeLibrary	Yes
HBA_GetNumberOfAdapters	Yes
HBA_GetAdapterName	Yes
HBA_OpenAdapter	Yes
HBA_CloseAdapter	Yes
HBA_GetAdapterAttributes	Yes
HBA_GetAdapterPortAttributes	Yes
HBA_GetDiscoveredPortAttributes	Yes
HBA_GetPortAttributesbyWWN	Yes
HBA_SendCTPassThru	Yes
HBA_SendCTPassThruV2	Yes
HBA_RefreshInformation	Yes
HBA_GetFcpTargetMapping	Yes
HBA_SendScsiInquiry	Yes
HBA_SendReportLuns	Yes
HBA_SendReadCapacity	Yes
HBA_GetPortStatistics	Yes
HBA_ResetStatistics	No
HBA_GetFcpPersistentBinding	No
HBA_GetEventBuffer	No
HBA_SetRNIDMgmtInfo	Yes
HBA_GetRNIDMgmtInfo	Yes
HBA_SendRNID	Yes
HBA_SendRNIDV2	Yes

TABLE 1-1 Supported and Unsupported SNIA Interfaces *(Continued)*

SNIA API	Sun StorEdge SAN Foundation Software Support
HBA_ScsiInquiryV2	Yes
HBA_ScsiReportLUNsV2	Yes
HBA_ScsiReadCapacityV2	Yes
HBA_OpenAdapterByWWN	Yes
HBA_RefreshAdapterConfiguration	Yes
HBA_GetVendorLibraryAttributes	Yes
HBA_GetFcpTargetMappingV2	Yes
HBA_SendRPL	No
HBA_SendRPS	No
HBA_SendSRL	No
HBA_SendLIRR	No
HBA_SendRLS	Yes
HBA_RemoveCallback	No
HBA_RegisterForAdapterEvents	No
HBA_RegisterForAdapterAddEvents	No
HBA_RegisterForAdapterPortEvents	No
HBA_RegisterForAdapterPortStatEvents	No
HBA_RegisterForTargetEvents	No
HBA_RegisterForAdapterLinkEvents	No
HBA_RegisterForAdapterTargetEvents	No
HBA_GetFC4Statistics	No
HBA_GetFCPStatistics	No
HBA_GetBindingCapability	No
HBA_GetBindingSupport	No
HBA_SetBindingSupport	No
HBA_SetPersistentBindingV2	No
HBA_GetPersistentBindingV2	No
HBA_RemovePersistentBinding	No
HBA_RemoveAllPersistentBindings	No

SNIA packages are now part of the Sun StorEdge SAN Foundation software stack. See the *Sun StorEdge SAN Foundation Software 4.4 Installation Guide* and *Sun StorEdge SAN Foundation Software 4.4 Release Notes* for more information on the location of the packages and how to install them. The package names are:

SNIA common library	SUNWcfclr
	SUNWcfcl
	SUNWcfclx
SAN Foundation Vendor Library	SUNWfchbr
	SUNWfchba
	SUNWfchbx
FCSM driver	SUNWfcsm
	SUNWfcsmx
SAN Foundation software	SUNwsan
	SUNWmdiu
	SUNWcfpl
	SUNWcfplx

Zones and Ports

Understanding zoning and port usage is fundamental to understanding the use of configuration rules with the supported hardware. This section explains the use of zones and ports in preparation of the next section, which covers the configuration rules. Topics covered include:

- “Zone Types” on page 6
- “Port Types” on page 7

Zone Types

Zoning is a function of the switch that allows segregation of devices by ports or World Wide Names (WWNs). You can create zones for a variety of reasons, such as security, simplicity, performance, or dedication of resources. This software supports hard zones, segmented loop (SL) zones, name server (NS) zones, and industry-standard port-based and WWN-based NS zones. See your third-party vendor documentation for more information.

The two types of NS zones discussed most frequently in this documentation include:

- Port-based NS zones
- WWN-based NS zones

Name Server Zones

NS zones use fabric protocols to communicate with Fibre Channel devices. Each NS zone defines which ports or devices receive NS information. The Sun StorEdge T3 arrays with firmware level 1.18.02 or higher and Sun StorEdge T3+ arrays with firmware level 2.1.04 or higher support loop (TL) port connections.

The Sun StorEdge T3+ array with firmware level 2.1.04 or higher supports fabric connections. FL ports are supported only for Sun StorEdge L180/L700 tape libraries. Please refer to your switch documentation for more information.

Segmented Loop Zones

The SFS 4.0 and higher releases do not support Segmented Loop (SL) zones or ports. If you have a SAN that requires SL connectivity to private loop devices, you can connect a new switch from the SFS 4.4 release to the same host as an old switch from previous releases. The two SANs must, however, each have separate directories on the host for their respective management tools.

Port Types

This table applies to Sun switches only.

TABLE 1-2 Sun Switch Port Types

Port Type	Description	Supported Devices
TL Ports	Translated loop	Storage devices connected to the Sun switch only.
FL Ports	Public loop	Sun StorEdge L180/L700 tape libraries.
F Ports	Point-to-point fabric	Host bus adapters, storage devices.
E Ports	Inter-switch port	Cascaded switches acting as ISLs, which are configured initially in fabric port mode.
G Ports	General ports	Automatically configure to E or F ports to support switches or fabric devices. All switch ports should be set to G-port, except for tape libraries that do not support F-port; see GL ports below.
Gx Ports	Public loop or general ports	Automatically configure to FL or G ports to support hosts or switches (Sun StorEdge Network 2 Gbit McDATA Sphereon 4300 and 4500 switches).
GL Ports	General loop ports	Automatically configure to FL, E, or F ports to support public loop, point-to-point, or switch devices. This port type is used only for setting L180/L700 tape libraries to FL.

Guidelines

This section covers guidelines for the following topics:

- “Cascading Guidelines” on page 8
- “IPFC Guidelines” on page 9
- “Multipathing Guidelines” on page 10

Cascading Guidelines

- Hub-to-switch connectivity is not supported in a fabric.
- Cascading requires connection of E ports from switch to switch with either a shortwave or longwave GigaBit Interface Converter (GBIC) or Small Form Factor Pluggable (SFP) 2-Gbit transceiver. The use of longwave SFPs and long haul fiber optics allows users to reach geographically separated storage and servers, perhaps for disaster recovery purposes.
- You can have a maximum of eight switches cascaded in a linear series. In other words, you can have seven inter-switch links between switches. Two of the ISLs can be long-wave connections. Check with your switch vendor for specific ISL count limitations because not all switches support a cascade of eight.
- Long-wave and short-wave ISLs can be used to cascade switches. The maximum distance for a single ISL hop is 10 km. Check with your third-party vendor documentation for details.
- If 1 Gbit and 2 Gbit switches are used together, a maximum of 16 switches can be connected together.
- If only 2 Gbit switches are used, a maximum of 24 switches can be connected together. Check with your switch vendor for specific switch count limitations because some switches might support a different number of switches or combinations.
- Any number of ISLs can be used between two switches.

IPFC Guidelines

IP over Fibre Channel (IPFC) devices are supported for use with Network File System (NFS) software, Network Attached Storage (NAS) devices and Sun StorEdge Network Data Replicator (Sun SNDR) software, or Sun StorEdge Availability Suite 3.1 remote mirror software.

TABLE 1-3 IPFC (NFS/NAS and SNDR)

Feature	Supported
Cascading	Yes, with Fabric NS zones only
Zone type	Fabric NS zone (with the HBA configured as an F port point-to-point connection)
Maximum number of device ports per zone	4*

* With the Sun StorEdge PCI Dual Fibre Channel Network Adapter+, only physical port 2 can be used for IPFC. With the Sun StorEdge cPCI Dual Fibre Channel Network Adapter, only physical port 1 can be used for IPFC.

The following restrictions apply:

- IPFC is supported for 2 Gbit Sun switches only; it is not supported on 1 Gbit Sun switches.
- Promiscuous mode is not supported. The `snoop (1M)` utility cannot be used.
- Multicasting is supported through broadcasting only.
- Assign the IP address of the IPFC port to a subnet different from that of the Ethernets on the same system.
- Network cards using IPFC can not be used as routers. The `/etc/notrouter` file must be present on the host.
- When using IPFC, storage devices and hosts should be in separate zones. The storage device should have one path to one zone and another path to another zone for failover and redundancy. The host can have more than one path to a specified zone, and it should have at least one path to each zone so that it can see the respective storage.
- Failover is not possible with IPFC or FC tape.

Multipathing Guidelines

- Before you configure port-based zones to storage devices, you might want to configure volumes, or LUNs, for those devices.
- If you want high availability, you should also enable the multipathing capability for load balancing of I/O traffic between hosts and arrays in the SAN.

Multipathing capability is provided by the Sun StorEdge Traffic Manager driver which runs in the Solaris operating environment. See the *Sun StorEdge Traffic Manager Software Installation and Configuration Guide* for details about the component. If you are using another multipathing application, see the documentation for that application.

Note – The Sun StorEdge Traffic Manager software features are not available for SCSI devices but are available for Fibre Channel target devices. Multipathing is not supported on tape drives or tape libraries.

Configuring the Switch Ethernet Port

After you install the SFS, you must configure the Sun switch Ethernet port. This chapter describes the following topics:

- “Using RARP to Change the IP Address” on page 12
- “To Prepare the File System for RARP” on page 12
- “To Set the Network IP Address With RARP” on page 14

Note – The following RARP instructions apply to Sun switches only. Other switch setup procedures can be found in the documentation for your switch.

The IP address of Sun switches is set to the default value of 10.0.0.1. You can change this default IP address by using either the Reverse Address Resolution Protocol (RARP) or the switch management software supplied with your switch.

This guide provides the procedures for changing the IP address using only RARP. See the user manual provided with your switch hardware for further information on configuring the switch.

Using RARP to Change the IP Address

You can use RARP to change the address when you first receive the switch, or if you forget the address, at a later time.

The RARP feature automatically downloads the assigned IP address to the switch. After you cable the switch and turn the power on, the IP address is automatically assigned.

You will need the following information to change the IP address:

IP address _____	Contact the person who maintains your network for this information.
MAC address _____	See the documentation that came with your switch for the location of the MAC address.
Switch _____	This is the assigned name of the switch that you are installing.

▼ To Prepare the File System for RARP

1. Verify that the boot method on the switch is set to RARP.

Check the documentation that comes with the hardware for instructions on how to verify that the boot method is set to RARP, which the default set in the factory.

2. On a host connected to the same subnet as the switch, save your current configuration.

3. Check that the `/etc/ethers` file exists.

```
# ls /etc/ethers
```

If the `ethers` file does not exist, create one.

4. If you find the file, create the `/etc/ethers.SAV` backup file:

```
# cp -p /etc/ethers /etc/ethers.SAV
```

5. Back up the affected management station configuration files:

```
# cp -p /etc/nsswitch.conf /etc/nsswitch.conf.SAV
# cp -p /etc/hosts /etc/hosts.SAV
```

6. On a host connected to the same subnet as the switch, edit the `/etc/ethers` file by adding the MAC address and switch name.

For example:

```
8:0:20:7d:93:7e switch-name
```

In this example:

- `8:0:20:7d:93:7e` is the MAC address
- `switch-name` is the name of the switch you are installing.

7. Edit the `/etc/hosts` file with the IP address and switch name.

For example:

```
192.129.122.111 switch-name
```

In this example, `192.129.122.111` is the assigned IP address.

8. Edit the `/etc/nsswitch.conf` file to reference the local system files.

To ensure that the Solaris operating environment uses the changes made to `/etc/ethers` and `/etc/hosts` files, edit the `host` and `ethers` entries in the `/etc/nsswitch.conf` file so that the `files` parameter appears before the `[NOTFOUND=return]` statements.

```
hosts:nis files [NOTFOUND=return]
ethers:nis files [NOTFOUND=return]
```

▼ To Set the Network IP Address With RARP

1. Determine whether the RARP daemon is running.

```
# ps -eaf | grep rarpd
```

2. If RARP is not running, start the RARP daemon in the Solaris software environment.

```
# /usr/sbin/in.rarpd -ad
```

The IP address automatically downloads to the switch after you install the switch and power on the system.

Note – RARP might not work properly if you have connected the Ethernet port of the switch to an Ethernet switch. This is due to the long time it might take for the Ethernet switch to transmit the RARP. In such a case, you can use the management tools to change the IP address. See the documents that came with your switch.

3. Power off the switch and then power it back on.
4. Verify the switch IP address.

```
# ping ip-address
```

For example:

```
# ping 192.129.122.111
```

5. If the RARP daemon has started, return to the window in which the daemon was activated and press the Control-C key sequence to stop the daemon operation.
If the IP address is not received from the RARP server, it retries three times before reverting to the original IP address (factory default setting).
6. To return the host to its original state copy the `nsswitch.conf`, `ethers`, and `hosts` files that you saved in “To Prepare the File System for RARP” on page 12.
7. Make the IP address setting permanent.

The IP address setting for the switch that you obtain using RARP is not a permanent setting until RARP is disabled and the switch is reset or power cycled.

Booting From Fibre Channel Storage Devices

Booting from storage devices on a network through a host bus adapter (HBA) is mandatory when the host is an enterprise-level system or other system that does not have its own directly connected disks. It is optional and preferable when a host has a directly connected disk.

This chapter contains two procedures for booting from the network adapter to enable detection of the installation devices. You must boot from the adapter under these conditions:

- The network adapter card driver is not bundled with the Solaris OE on the boot server network.
- The host is a system that does not have disks connected directly to it, such as an enterprise-level system.

Note – Only an experienced UNIX system administrator should perform the boot procedures in this chapter.

This chapter covers the following topics:

- “Special Bootability Considerations” on page 16
- “Determining Which Boot Method to Use” on page 17
- “Net Install Patching Method” on page 18
- “Dump and Restore Method” on page 24

Special Bootability Considerations

When you decide to boot from a device on the network, remember the following considerations.

- Maximize access to the boot device by:
 - Keeping Solaris boot images on a LUN that is not overly subscribed on the device level.
 - Keeping applications and additional software on separate devices from the boot device and boot path.
 - Reducing the physical proximity of host and fabric boot devices as much as possible, and definitely keep them less than 10 km apart.

If the HBA, path, storage array, or physical device that contains or provides access to the boot device is busy, removed, blocked, or made unavailable to the host for a period of time the host may panic.

- Use the procedures in this chapter if you are not using a CD 0 and if the HBA driver is not bundled with the Solaris operating environment on the boot server on the network. Check your Solaris operating environment to determine whether the driver is bundled with it.
- Remove boot disk from Logical Manager control before you begin the fabric boot procedure.
- Before performing fabric boot, execute the `luxadm qlgc_s_download` or `luxadm fcode_download` (for JNI HBAs) commands to list all the current FCode levels. You must have the most recent FCode level installed on the HBA. See “Software Installation Sequence” in the *Sun StorEdge SAN Foundation Software 4.4 Installation Guide* for the base patch list. Follow the instructions on the README distributed with the patch.
- Bootability on fabric devices is only possible with Sun StorEdge 1 Gbit and 2 Gbit FC PCI Single and Dual Channel Network adapters and Sun StorEdge JNI 2 Gbit PCI Single and Dual Channel Network adapters. See the *Sun StorEdge SAN Foundation Software 4.4 Release Notes* for additional information on device firmware requirements.)
- Starting with Sun StorEdge SAN Foundation Software version 4.4, you can boot from a Sun StorEdge Traffic Manager (STMS)-enabled device.

If you are using versions earlier than 4.4, you cannot boot from an STMS-enabled disk link or port. Disable Sun StorEdge Traffic Manager on any HBA port, or path used for fabric boot. To remove a path from Traffic Manager control, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*.

An STMS-enabled disk link looks like:

```
/dev/rdisk/c18t60020F200000457C3A83208A000825FAd0s0
```


If Traffic Manager is not disabled, it is possible to create a boot image that will not be accessible to the host upon the completion of the procedures in this section.

- Use array-based LUN masking or host-based configuration management with `cfgadm(1M)` to allow only devices or LUNs required for booting on the path used for fabric boot.
- Perform the “Dump and Restore Method” procedures while the system is booted in single user mode. This will maintain file system consistency during the placement of the boot image on the fabric device and eliminate the risk of the host panicking due to device inaccessibility.
- Return the boot disk to be under Logical Manager control if required. Create a mirror of the fabric boot disk to another disk if required.

Determining Which Boot Method to Use

You can boot from the small kernel image on the adapter in two ways:

- Using a boot or installation server (or the net install patching method)
- Using a temporarily connected boot disk (or the dump and restore method)

The net install patching method requires a JumpStart server and network connectivity, and it is more suitable for large system configurations than for small ones. This method is easier than the Dump and Restore method. The net install patching method supports non-fabric devices only. Do not use this method to boot from fabric devices. When the system is using an SBus HBA, use the net install patching method only. Otherwise, the bootability features are already part of the OE and the steps are not needed.

The dump and restore method is more difficult than the net install patching method. However, the method is available to anyone with existing systems, whether small or large configurations, and it provides a disk image when you have completed the process. You can boot from either fabric or non-fabric devices with the Dump and Restore method.

Net Install Patching Method

The procedures in this section explain how to install the Solaris OE from a boot or install server onto the host with the host adapter. Any other host on the same subnet can be set up as a boot or install server. The Net Install method is not to be used with a fabric boot device; it must be used with direct-connect storage or private loop attached storage *only*.

Note – The procedure is identical whether you are using a boot server or an install server.¹ Therefore, the convention used in this chapter is to refer to either type of server as the “boot or install server.”

The topics in this section include:

- “Overview Of the Net Install Method” on page 19
- “To Set Up the Boot or Install Server” on page 20
- “To Modify the Boot Image” on page 21
- “To Set Up the Client and Reboot” on page 22

1. For the distinction between these two types of servers, see the Solaris system administration documentation.

Overview Of the Net Install Method

Installing a client from a boot or install server uses two images of the Solaris OE:

- A boot mini-root (which is referred to from here on as the *boot image*)
- A separate *install image* that gets copied onto the boot disk

This procedure assumes you know how to install Solaris software over the network as described in the Solaris installation manuals. For more information, see the `man(1)` pages for the commands that are used in this procedure.

Note – While you are setting up the boot or install server, you must either have the contents of the Solaris installation CD-ROM copied to a disk that is connected to the boot or install server, or have the Solaris installation CD inserted and mounted from an attached CD-ROM device.

You must enable bootability using a boot or install server to make both images aware of the host adapter. At the start of the procedure, download the driver and all required patches to an exported directory on the boot or install server.

For actual instructions, begin with “To Set Up the Boot or Install Server” on page 20.

Note – You cannot use JumpStart until you modify the install image.

- “To Set Up the Boot or Install Server” on page 20
- “To Modify the Boot Image” on page 21
- “To Set Up the Client and Reboot” on page 22

▼ To Set Up the Boot or Install Server

1. **Become superuser on the host to be used as the boot or install server.**
2. **Use the `setup_install_server(1M)` command from the `Tools` directory in the location where the Solaris software resides.**

As shown in the following screen example, the `setup_install_server` command copies the boot image to a directory on the boot or install server. (The boot directory is named `/original_OS_dir/Boot` in the example.) The example shows the command being run from the `Tools` directory on a mounted Solaris 8 installation CD.

Note – This step might take 20 minutes or more to complete.

```
# cd /cdrom/cdrom0/s0/Solaris_8/Tools
# ./setup_install_server -t /original_OS_dir/Boot /new_OS_copy_dir
```

3. **Download the driver packages and patches into the `/export_public` directory on the boot or install server.**

See *Sun StorEdge SAN Foundation Software 4.4 Installation Guide* for installation details. If the `/export_public` directory already exists, create another directory with a new name and use that in place of `/export_public` in the rest of the instructions.



Caution – Do not use the `pkgadd(1M)` and `patchadd(1M)` commands in “To Manually Install the SFS Packages and Patches” in the *Sun StorEdge SAN Foundation Software 4.4 Installation Guide* to install the software because this would modify the install server's own installation.. Rather, see “To Modify the Boot Image” on page 21 to install relocatable versions of the packages.

▼ To Modify the Boot Image

1. Install the driver packages into the boot image.

The following example installs all packages previously downloaded into the `/export_public` directory where *Solaris_version* is the version of the OE you are installing.

Note – See the *Sun StorEdge SAN Foundation Software 4.4 Installation Guide* for details to install the packages and patches. You must follow the specified installation order for a successful installation so the driver will run.

```
# cd /export_public
# pkgadd -R /new_OS_copy_dir/Solaris_version/Tools/Boot -d . package_name
```

2. Install all required patches into the boot image.

The following example installs any patch(es) that were previously downloaded into the `/export_public` directory. Repeat the `patchadd` command for all the patches you need to add.

```
# cd /export_public
# patchadd -C /new_OS_copy_dir/Solaris_version/Tools/Boot -M /export_public patch_ID
```

3. Make sure the host name of the host getting the new OE, its IP address, and its Ethernet address have been added to the name service (`/etc` files, NIS, or NIS+).

4. Run the `add_install_client(1M)` command to add the host with the host adapter as a boot or install client.

The command is in the `/export_public/new_OS_copy_dir/Solaris_version/Tools`.

```
# add_install_client host_name platform_name
```

Note – You can find the platform name by running the `uname` command with the `-m` option on the host that has the host adapter.

5. Log out of the boot or install server.

▼ To Set Up the Client and Reboot

1. Bring the client host (with the HBA) down to the `ok` prompt at run level 0.

See the Solaris system administration documentation for the commands that can be used with different configurations. The following example uses the `shutdown(1M)` command.

```
# shutdown -i 0
...
ok
```



Caution – Do not reboot the boot or install server while installing the new OE on the client.

2. Boot the host from the net.

```
ok boot net
```

The Solaris interactive installation program runs from the boot or install server.

3. Respond to the prompts according to your configuration as instructed in the Solaris installation guide.

Specify the new boot disk as the destination for the operating environment installation.

4. When prompted to choose between automatic reboot or manual reboot, click the Manual Reboot button, complete the remaining question, and then start the installation.

This question offering a choice between automatic and manual reboot is the last question before the installation starts. If you are using the `suninstall` program, choose `boot manual`.

5. Mount the `/export_public` directory, which contains the driver packages and any needed patches, onto the `/mnt` directory mount point.

The following example uses `boot_install_server` as the name of the boot or install server.

```
# mount boot_install_server:/export_public /a/mnt
```

6. Install the driver packages into the install image.

The following example installs all packages previously downloaded in “To Set Up the Boot or Install Server” on page 20 into the `/export_public` directory. When prompted, install the packages in this order: SUNWsan, SUNWcfpl, SUNWcfplx.

```
# cd /a/mnt
# pkgadd -R /a -d . package_name
```

7. Install any needed patches into the boot image.

The following example installs all patches that were previously downloaded into the `/export_public` directory.

Note – Install the patches in their appropriate order to ensure the installation is successful.

```
# cd /a/mnt
# patchadd -R /a patch_ID
```

8. Bring the system down to the `ok` prompt at run level 0.

```
# halt
```

9. Reboot the host from the newly installed operating environment.

```
ok boot -r
```

Dump and Restore Method

Note – These examples use disk 0 as the temporary disk (c0t1d0) and disk 5 (c8t50060E80034E7818d0) as the new boot disk. If you were selecting a direct attached device as the new boot disk, you would select a link path that might look like: /dev/rdisk/c7t16d0 instead of /dev/rdisk/c8t50060E80034E7818d0. You should not select links that look like:
: /dev/rdisk/c4t50020F23000043D9d0s0 instead of /dev/rdisk/c7t1d0. You should not select links that look like:
/dev/rdisk/c18t60020F200000457C3A83208A000825FAd0s0, which are STMS-enabled devices. It is not possible to install a boot image on such a path.

The topics in this section include:

- “Overview Of the Dump and Restore Method” on page 24
 - “Partitioning the New Boot Disk the Same As the Temporary Boot Disk” on page 25
 - “Creating the New Boot Files” on page 35
-

Note – The dump and restore procedure must be done in single-user mode to maintain file system consistency during the placement of the boot image on the fabric device; otherwise, the boot image may not be fully accessible, the host may panic, and the procedures in this section will need to be performed again.

Overview Of the Dump and Restore Method

You can use the dump and restore method to boot from both fabric and non-fabric devices. To enable bootability using a temporary boot disk, a boot disk must be directly connected, at least temporarily, to the host. The boot disk must have the following installed:

- Solaris OE
 - SFS packages and any needed patches
-

Note – You can remove the initial boot disk if it is not needed after the boot disk is enabled.

Partitioning the New Boot Disk the Same As the Temporary Boot Disk

This section contains several sub procedures you must perform to complete the first phase of enabling bootability from a temporarily connected boot disk. These sub-procedures include:

- “To Prepare To Partition the New Disk” on page 25
- “To Record the Partition Layout” on page 26
- “To Change to the New Boot Disk” on page 29
- “To Specify Slices on the New Boot Disk” on page 31
- “To Label the New Boot Disk” on page 33
- “To Create File Systems on the New Boot Disk” on page 34

▼ To Prepare To Partition the New Disk

1. **Disable the Sun StorEdge Traffic Manager (STMS) software globally or for each pHCI.**

If you choose to disable STMS on each pHCI, make sure that it is the HBA through which the fabric boot device is seen.

2. **Become superuser on the host with the host adapter.**
3. **If the driver packages and patches are not already installed, download the software and install it on the host.**

See the *Sun StorEdge SAN Foundation Software 4.4 Installation Guide* for details.



Caution – Do not use the `pkgadd(1M)` and `patchadd(1M)` commands in “To Manually Install the SFS Packages and Patches” in the *Sun StorEdge SAN Foundation Software 4.4 Installation Guide* to install the software. Rather, see “To Modify the Boot Image” on page 21 to install relocatable versions of the packages.

4. **Shut down and restart your system.**

```
# /touch reconfigure
# shutdown -y -g0 -i 6
```

5. **Log in to the host as root.**

▼ To Record the Partition Layout

After you log back into the host, you can record the layout of the partitions, or slices, on the system boot disk.

1. Enter the `format(1M)` command so that the operating environment recognizes the devices attached to the card.

If needed, see the `format` man page and the instructions about adding a disk and using the `format` command in the Solaris administration documentation.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
 0. c0t1d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
    /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w2100002037e43542,0
 1. c0t2d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
    /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000020374205a1,0
 2. c7t16d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
    /pci@8,600000/SUNW,qlc@1/fp@0,0/ssd@w500000e01002b7c1,0
 3. c7t17d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
    /pci@8,600000/SUNW,qlc@1/fp@0,0/ssd@w500000e0100357c1,0
 4. c7t18d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
    /pci@8,600000/SUNW,qlc@1/fp@0,0/ssd@w500000e0100369e1,0

Specify disk (enter its number):
```

2. Specify the disk where the operating system is installed on the temporary boot disk.

The following screen example specifies disk 0.

```
Specify disk (enter its number): 0
```

3. Use the partition command to display the PARTITION MENU.

```
format> partition
PARTITION MENU:
  0      - change '0' partition
  1      - change '1' partition
  2      - change '2' partition
  3      - change '3' partition
  4      - change '4' partition
  5      - change '5' partition
  6      - change '6' partition
  7      - change '7' partition
select  - select a predefined table
modify  - modify a predefined partition table
name    - name the current table
print   - display the current table
label   - write partition map and label to the disk
!<cmd> - execute <cmd>, then return
quit
partition>
```

4. Use the print command to display the partition table for the specified disk.

```
partition> print
Current partition table (original):
Total disk cylinders available: 7506 + 2 (reserved cylinders)

Part      Tag      Flag      Cylinders      Size      Blocks
  0      root      wm        0 - 55      128.84MB  (56/0/0)  263872
  1      swap      wu        56 - 111    128.84MB  (56/0/0)  263872
  2      backup    wu        0 - 7505    16.86GB   (7506/0/0) 35368272
  3 unassigned  wm         0           0         (0/0/0)    0
  4 unassigned  wm         0           0         (0/0/0)    0
  5 unassigned  wm         0           0         (0/0/0)    0
  6      usr      wm       112 - 7505    16.61GB   (7394/0/0) 34840528
  7 unassigned  wm         0           0         (0/0/0)    0
```

5. Record the layout (sizes and numbers) assigned to the slices on the temporary boot disk, and then type `quit`.

The `FORMAT MENU` is displayed.

```
partition> quit
FORMAT MENU:
  disk      - select a disk
  type      - select (define) a disk type
  partition - select (define) a partition table
  current   - describe the current disk
  format    - format and analyze the disk
  repair    - repair a defective sector
  label     - write label to the disk
  analyze   - surface analysis
  defect    - defect list management
  backup    - search for backup labels
  verify    - read and display labels
  save      - save new disk/partition definitions
  inquiry   - show vendor, product and revision
  volname   - set 8-character volume name
  !<cmd>    - execute <cmd>, then return
  quit

format>
```

▼ To Change to the New Boot Disk

After you record the partition layout, change to the new boot disk.

1. **On the target controller containing the destination boot image, configure the fabric device to create device nodes on each server connected to the SAN:**

```
# cfgadm -c configure ap_id
```

2. **Type `format` to list the disks available to change the current disk to the new boot disk:**

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
 0. c0t1d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
    /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w2100002037e43542,0
 1. c0t2d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
    /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000020374205a1,0
 2. c7t16d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
    /pci@8,600000/SUNW,qlc@1/fp@0,0/ssd@w500000e01002b7c1,0
 3. c7t17d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
    /pci@8,600000/SUNW,qlc@1/fp@0,0/ssd@w500000e0100357c1,0
 4. c7t18d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
    /pci@8,600000/SUNW,qlc@1/fp@0,0/ssd@w500000e010032331,0
 5. c8t50060E80034E7818d0 <HITACHI-OPEN-L-SUN-2101 cyl 49437
    alt 2 hd 15 sec 96
    pci@8,600000/SUNW,qlc@2/fp@0,0/ssd@w50060e0034e7818,0"

Specify disk (enter its number):
```

3. **Make a note of the device path name of the new boot disk.**

For example, for disk 5, the new boot disk in this example, the device path name shown is: `/pci@8,600000/SUNW,qlc@2/fp@0,0`.

Note – The `qlc` number changes, depending on the type of host and slot in use.

4. Type the number of the disk to be formatted to display the FORMAT menu.

The following screen example uses disk 5.

```
Specify disk (enter its number): 5
selecting c8t50060E80034E7818d0:
[disk formatted]

FORMAT MENU:
  disk      - select a disk
  type     - select (define) a disk type
  partition - select (define) a partition table
  current  - describe the current disk
  format   - format and analyze the disk
  repair   - repair a defective sector
  label    - write label to the disk
  analyze  - surface analysis
  defect   - defect list management
  backup   - search for backup labels
  verify   - read and display labels
  save     - save new disk/partition definitions
  inquiry  - show vendor, product and revision
  volname  - set 8-character volume name
  !<cmd>   - execute <cmd>, then return
  quit

format>
```

5. Make a note of the device name of the disk.

The device name of the disk in the previous screen example is
c8t50060E80034E7818d0.

▼ To Specify Slices on the New Boot Disk

After you change to the new boot disk, specify one slice on the new boot disk for every slice on the temporary boot disk. The following examples specify the root slice 0 on the new boot disk to match slice 0 on the temporary boot disk.

1. Use the `partition` command to display the PARTITION MENU.

```
format> partition
PARTITION MENU:
    0      - change '0' partition
    1      - change '1' partition
    2      - change '2' partition
    3      - change '3' partition
    4      - change '4' partition
    5      - change '5' partition
    6      - change '6' partition
    7      - change '7' partition
select    - select a predefined table
modify    - modify a predefined partition table
name      - name the current table
print     - display the current table
label     - write partition map and label to the disk
!<cmd>   - execute <cmd>, then return
quit
partition>
```

2. Type the number of the slice to be defined (in this case, slice 0).
3. Enter the partition ID tag.

The following example shows a question mark (?) entered after the prompt. The list of accepted partition ID tags is displayed. You can accept the default partition ID tag of root by pressing the Return key.

```
Enter partition id tag[root]: ?
Expecting one of the following: (abbreviations ok):
    unassigned    boot        root        swap
    usr           backup     stand       var
Enter partition id tag[root]:
Enter partition permission flags[wm]:
```

4. Type the partition permission flags.

You can accept the default permission flags `wm` by pressing the Return key.

```
Enter partition permission flags[wm]:  
Enter new starting cyl[0]:
```

5. Type the new starting cylinder.

You can accept the default new starting cylinder of 0 accepted by pressing the Return key.

```
Enter new starting cyl[0]:  
Enter partition size[262899b, 91c, 128.37mb, 0.13gb]:
```

6. Type the partition size.

The destination disk must have partitions the same size or larger. For example, if the source disk is a 9 Gbyte disk and the destination disk is 146 Gbyte, you would lose 137 Gbyte by making them the same. The following example shows the partition size of 2.00gb entered.

```
Enter partition size[262899b, 91c, 128.37mb, 0.13gb]: 2.00gb  
partition>
```

7. Use the `print` command to display the updated partition table.

8. Repeat Step 2 through Step 7 as needed until all slices are defined as they are in the source boot disk.

9. Type `quit` to return to the `FORMAT` menu.

```
partition> quit

FORMAT MENU:
  disk      - select a disk
  type      - select (define) a disk type
  partition - select (define) a partition table
  current   - describe the current disk
  format    - format and analyze the disk
  repair    - repair a defective sector
  label     - write label to the disk
  analyze   - surface analysis
  defect    - defect list management
  backup    - search for backup labels
  verify    - read and display labels
  save      - save new disk/partition definitions
  inquiry   - show vendor, product and revision
  volname   - set 8-character volume name
  !<cmd>    - execute <cmd>, then return
  quit

format>
```

▼ To Label the New Boot Disk

After you specify the slices on the new boot disk, label the new boot disk with the new partition table.

1. Begin the labeling process.

```
format> label
```

2. Type `y` to continue labeling the disk.

```
Ready to label disk, continue? y
```

3. When the labeling is complete, type `q` to quit the `format` program.

```
format> quit
#
```

▼ To Create File Systems on the New Boot Disk

Use this procedure to create a file system on the new boot disk for every slice on the temporary boot disk.

Perform this step to create a file system on the new boot disk for the `/root` and `/home` directory slices. You should also do this step for any other directories you created, except for the `/backup` and `/swap` partitions.

1. **Create a file system on each slice (other than S2) on the disk using the `newfs(1M)` command.**

Note – Do not create a file system on slice 2, `/backup`, or `/swap`, or else it will erase data on all the other slices.

Enter the `newfs` command followed by the device name of the slice. In this example for a non-fabric boot device, the device name for slice 0 of disk

`c8t50060E80034E7818d0` is `/dev/rdisk/c8t50060E80034E7818d0s0`.

```
# newfs /dev/rdisk/c8t50060E80034E7818d0s0
newfs: construct a new file system /dev/rdisk/c8t50060E80034E7818d0s0: (y/n)? y
/dev/rdisk/c8t50060E80034E7818d0s0:
4194828 sectors in 1452 cylinders of 27 tracks, 107 sectors
    2048.3MB in 46 cyl groups (32 c/g, 45.14MB/g, 7488 i/g)
super-block backups (for fsck -F ufs -o b=#) at:
 32, 92592, 185152, 277712, 370272, 462832, 555392, 647952, 740512, 833072,
 925632, 1018192, 1110752, 1203312, 1295872, 1388432, 1480992, 1573552,
 1666112, 1758672, 1851232, 1943792, 2036352, 2128912, 2221472, 2314032,
 2406592, 2499152, 2591712, 2684272, 2776832, 2869392, 2958368, 3050928,
 3143488, 3236048, 3328608, 3421168, 3513728, 3606288, 3698848, 3791408,
 3883968, 3976528, 4069088, 4161648,
```

For more information about how to create file systems, see the Solaris system administration documentation. Also, see the `newfs(1M)` man page.

2. **Repeat the procedure for all destination disk partitions except slice 2, `/backup`, and `/swap`.**
3. **Proceed to “To Copy the Contents of Non-Root File Systems Onto the New Boot Disk” on page 38.**

Creating the New Boot Files

This section contains several procedures, including:

- “To Copy the Boot Block and Root File System Contents to the New Boot Disk” on page 35
- “To Update the `vfstab` File” on page 37
- “To Copy the Contents of Non-Root File Systems Onto the New Boot Disk” on page 38
- “To Specify the New Boot Disk as the Boot Device” on page 39

▼ To Copy the Boot Block and Root File System Contents to the New Boot Disk

1. Install the boot block on the `root (/)` file system of the new disk.

The following example uses the `installboot(1M)` command to install the boot block. The boot block resides in the `/usr/platform/platform_name/lib/fs/ufs/bootblk` directory. The non-fabric example shows invoking the `uname` command with the `-m` option between left single quotes on the command line to specify the platform name.

```
# /usr/sbin/installboot /usr/platform/'uname -i'/lib/fs/ufs/bootblk  
/dev/rdisk/c7t50060E80034E7818d0s0
```

For more information, see the instructions on how to install a boot block in the *System Administration Guide, Volume 1* for Solaris 8 OE or *System Administration Guide: Basic Administration* for Solaris 9 OE.

2. Mount the `root` file system from slice 0 of the new boot disk onto the `/mnt` mount point.

```
# mount /dev/dsk/c8t50060E80034E7818d0s0 /mnt
```

3. Use the `ufsdump (1M)` and `ufsrestore (1M)` commands to copy the contents of the `root` file system from the temporary boot disk to the root slice of the new boot disk (on the `/mnt` mount point).

Make sure you use the original root slice in the command. In the following non-fabric example, the root slice is `/dev/rdisk/c0t1d0s0`.

```
# ufsdump 0f - /dev/rdisk/c0t1d0s0 | ( cd /mnt; ufsrestore rf -)
DUMP: Writing 32 Kilobyte records
DUMP: Date of this level 0 dump: Tue 19 Feb 2002 02:44:35 PM PST
DUMP: Date of last level 0 dump: the epoch
DUMP: Dumping /dev/rdisk/c0t1d0s0 (hba2-81://) to standard output.
DUMP: Mapping (Pass I) [regular files]
DUMP: Mapping (Pass II) [directories]
DUMP: Estimated 1818082 blocks (887.74MB).
DUMP: Dumping (Pass III) [directories]
DUMP: Dumping (Pass IV) [regular files]
Warning: ./lost+found: File exists
DUMP: 88.77% done, finished in 0:01
DUMP: 1818046 blocks (887.72MB) on 1 volume at 1363 KB/sec
DUMP: DUMP IS DONE
```

▼ To Update the `vfstab` File

After you copy the boot block and root files, update the `vfstab` file.

1. Change directories to `/mnt/etc` and open the `vfstab(4)` file for editing.

The following example shows the file systems defined.

```
# cd /mnt/etc
# vi vfstab
...
/dev/dsk/c0t1d0s1      -      -      swap      -      no      -
/dev/dsk/c0t1d0s0      /dev/rdisk/c0t1d0s0      /      ufs      1      no -
/dev/dsk/c0t1d0s7      /dev/rdisk/c0t1d0s7      /home    ufs      2      yes -
```

2. Replace the name of the temporary boot disk with the name of the new boot disk, and then save and quit the file.

The following example shows the disk name `c0t0` changed to `c8t50060E80034E7818d0` in the mount table entries for slices 0, 1, and 7.

```
/dev/dsk/c8t50060E80034E7818d0s1      -      -      swap      -      no      -
/dev/dsk/c8t50060E80034E7818d0s0      /dev/rdisk/c8t50060E80034E7818d0s0 /
ufs      1 no      -
/dev/dsk/c8t50060E80034E7818d0s7      /dev/rdisk/c8t50060E80034E7818d0s7
/home    ufs      2 yes -
:wq
#
```

3. Change to the `root` directory and unmount the mount point:

```
# cd /
```

4. Unmount the `root` file system on slice 0 from the `/mnt` mount point.

```
# umount /mnt
```

▼ To Copy the Contents of Non-Root File Systems Onto the New Boot Disk

Use this procedure for the `/export/home` directory and any other directories you created.

1. Mount the file system onto the `/mnt` mount point.

This example shows the copying of the `/home` file system from slice 7 to the new boot disk.

```
# mount /dev/dsk/c8t50060E80034E7818d0s7 /mnt
```

2. Use the `ufsdump(1M)` and `ufsrestore(1M)` commands to copy the contents of the file system from the temporary boot disk to the new boot disk.

```
# ufsdump 0f - /dev/rdisk/c0t1d0s7 | ( cd /mnt; ufsrestore rf -)
DUMP: Writing 32 Kilobyte records
DUMP: Date of this level 0 dump: Tue 19 Feb 2002 02:44:35 PM PST
DUMP: Date of last level 0 dump: the epoch
DUMP: Dumping /dev/rdisk/c0t1d0s7 (hba2-81://) to standard output.
DUMP: Mapping (Pass I) [regular files]
DUMP: Mapping (Pass II) [directories]
DUMP: Estimated 1818082 blocks (887.74MB).
DUMP: Dumping (Pass III) [directories]
DUMP: Dumping (Pass IV) [regular files]
Warning: ./lost+found: File exists
DUMP: 88.77% done, finished in 0:01
DUMP: 1818046 blocks (887.72MB) on 1 volume at 1363 KB/sec
DUMP: DUMP IS DONE
#
```

3. Unmount the file system from the `/mnt` mount point.

```
# umount /mnt
```

4. Repeat Step 1 through Step 3 as needed until you have copied the contents of all the file systems required.

▼ To Specify the New Boot Disk as the Boot Device

In this example, you specify a new boot device, create a device alias for the new boot device, and then boot from the new device. (Creating an alias simplifies the ability to boot from the device in the last step.)

1. Bring the host with the host adapter down to the `ok` prompt at run level 0.

See the information about shutting down a host in the Solaris system administration documentation for a list of the commands that can be used with different configurations. The following screen example uses the `shutdown(1M)` command.

```
# shutdown -y -g 10 -i 0
...
ok
```

2. At the `ok` prompt, type the following command to prevent the system from hanging:

```
ok reset-all
Resetting ...
.
.
.
ok
```

3. Make sure you can display your fabric or fc_al (Fibre Channel arbitrated loop) devices from the ok prompt. Use the show-devs command to list the available devices.

In this case, for fabric devices with Sun single and dual channel Fibre Channel PCI HBAs, the HBA firmware revision level must be 1.14.01 or higher.

```
ok show-devs
.
.
.
/pci@1f,2000/SUNW,qlc@1,1
/pci@1f,2000/SUNW,qlc@1
/pci@1f,2000/SUNW,qlc@1,1/fp@0,0
/pci@1f,2000/SUNW,qlc@1,1/fp@0,0/disk
/pci@1f,2000/SUNW,qlc@1/fp@0,0
/pci@1f,2000/SUNW,qlc@1/fp@0,0/disk
/pci@1f,4000/SUNW,qlc@5
/pci@1f,4000/SUNW,qlc@4
/pci@1f,4000/SUNW,qlc@2,1
/pci@1f,4000/SUNW,qlc@2
/pci@1f,4000/scsi@3
/pci@1f,4000/network@1,1
/pci@1f,4000/ebus@1
/pci@1f,4000/SUNW,qlc@5/fp@0,0
/pci@1f,4000/SUNW,qlc@5/fp@0,0/disk
/pci@1f,4000/SUNW,qlc@4/fp@0,0
/pci@1f,4000/SUNW,qlc@4/fp@0,0/disk
/pci@1f,4000/SUNW,qlc@2,1/fp@0,0
/pci@1f,4000/SUNW,qlc@2,1/fp@0,0/disk
/pci@1f,4000/SUNW,qlc@2/fp@0,0
/pci@1f,4000/SUNW,qlc@2/fp@0,0/disk
```

For this output example, the boot device is /pci@1f,4000/SUNW,qlc@4/. (You can determine the physical path of the device c#t#d#s# by using the format or luxadm display /dev/rdisk/c#t#d#s# commands.)

4. Use the **select** and **show-children** commands to display the devices attached to the specified path.

```
ok select /pci@1f,4000/SUNW,qlc@4/
ok show-children
***** Fabric Attached Devices *****
Adapter portId - 21300
Device PortId 21900 DeviceId 1 WWPN 50060E80034E7818
Lun 0 DISK HITACHI OPEN-9*6 -SUN2104
Lun 1 DISK HITACHI OPEN-9*6 -SUN2104
Lun 2 DISK HITACHI OPEN-V -SUN2104
Lun 3 DISK HITACHI OPEN-V -SUN2104
```

In this example, LUN 0 is the boot device.

5. Create a device alias for the boot device.

User input is in **bold** text. Use Control-Y to insert the entire device name at the **ok** prompt as instructed by the **show-disks** command.

Note – When you create the boot device for a fabric or fc_al disk device, you must use the device listed by the **show-disks** command. Note that in the format or luxadm command output, the device has **ssd@** and not **disk@** in the physical path. You must use the **disk@** portion; otherwise, you will not be able to boot from the device.

```
ok show-disks
a) /pci@1f,2000/SUNW,qlc@1,1/fp@0,0/disk
b) /pci@1f,2000/SUNW,qlc@1/fp@0,0/disk
c) /pci@1f,4000/SUNW,qlc@5/fp@0,0/disk
d) /pci@1f,4000/SUNW,qlc@4/fp@0,0/disk
e) /pci@1f,4000/SUNW,qlc@2,1/fp@0,0/disk
f) /pci@1f,4000/SUNW,qlc@2/fp@0,0/disk
g) /pci@1f,4000/scsi@3/disk
h) /pci@1f,4000/ebus@1/fdthree@14,3023f0
q) NO SELECTION
Enter Selection, q to quit: d
/pci@1f,4000/SUNW,qlc@4/fp@0,0/disk has been selected.
Type ^Y ( Control-Y ) to insert it in the command line.
e.g. ok nvalias mydev ^Y for creating devalias mydev for
/pci@1f,4000/SUNW,qlc@4/fp@0,0/disk
ok nvalias hds-fabric-boot
/pci@1f,4000/SUNW,qlc@4/fp@0,0/disk@w50060E80034E7818,0:a
```

6. Verify that the device alias `hds-fabric-boot` has been created.

```
ok devalias
screen                               /SUNW,ffb@1e,0
hds-fabric-boot
/pci@1f,4000/SUNW,qlc@4/fp@0,0/disk@w50060E80034E7818,0:a
.
.
.
```

7. Boot from the device.

```
ok boot hds-fabric-boot
Boot device:
/pci@1f,4000/SUNW,qlc@4/fp@0,0/disk@w50060E80034E7818,0:a
File and args:
SunOS Release 5.9 Version Generic_112233-05 64-bit
Copyright 1983-2002 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
```

Note – If an error message such as `cannot open boot device` is displayed, ensure that you have not missed a step in this procedure.

If you want this device to be a Sun StorEdge Traffic Manager-enabled device, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*.

Adding and Removing SAN Devices

By now, you should have identified features and guidelines for connecting devices and hosts to the switches. This chapter helps you with setting up the devices in the zones on your switches. If you have not already done so, configure your ports and zones according to the instructions in your vendor-specific documentation.

This section covers addition of the supported storage devices, such as the Sun StorEdge T3+ array and Sun StorEdge 39x0, 69x0, and 99x0 storage arrays, to a SAN. Topics include:

- “Special Considerations” on page 44
- “Adding and Removing Devices” on page 45

Special Considerations

This section covers:

- “Port Choices” on page 44
- “Multipathing” on page 45

Port Choices

When configuring an array, the host port is connected to an F port and the array is connected to an F or TL port on the switch. The TL (translation loop) port, represents eight-bit addressing devices as 24-bit addressing devices and vice versa.

TABLE 4-1 Port Types and For Storage Devices

Port Type	Storage Device
loop*	Sun StorEdge T3 array
loop or fabric†	Sun StorEdge T3+ array
	Sun StorEdge 39x0 array
fabric	Sun StorEdge 69x0 array
	Sun StorEdge 6x20 array
loop or fabric	Sun StorEdge 99x0 array
loop	STK 9840b tape drive
loop or fabric	STK 9940b tape drive
loop	STK 9840 tape drive
loop, public loop‡	Sun StorEdge L180/L700 tape libraries

* TL port - Sun switches only

† F port

‡ FL port

Although you may connect a Sun StorEdge T3 array with a TL port, the host bus adapter recognizes it as a fabric device. Sun StorEdge T3+ arrays and the Sun StorEdge 3510FC, 39x0, 69x0, and 99x0 storage arrays should be connected with F ports as a 24-bit addressing device for fabric connectivity. The STK 9840B tape drive require F ports when connected to 2 Gbit switches. Sun StorEdge L5500/L6000 libraries are not connected to the SAN.

Multipathing

You can connect arrays and other storage devices to the SAN with or without multipathing capability. The multipathing driver discussed in this book is the Sun StorEdge Traffic Manager software (STMS) and runs on the Solaris operating environment. The Sun StorEdge Traffic Manager software (STMS) is an associated application that is bundled with the SFS. For detailed information about the STMS or other multipathing utilities, see the *Sun StorEdge Traffic Manager Software Installation and Configuration Guide* or the third-party multipathing software manuals.

Adding and Removing Devices

This section covers:

- “Adding a Storage Device” on page 45
- “Removing a Storage Device” on page 48

Adding a Storage Device

- If the STMS is not enabled, read the section “Creating and Removing Individual Device Nodes Without Multipathing Enabled” on page 52. See “Creating and Removing Multiple Device Nodes Without Multipathing Enabled” on page 61 for specific information.
- If the STMS is enabled, read the section “Creating and Removing Individual Device Nodes With Multipathing Enabled” on page 65. See “Creating and Removing Multiple Device Nodes With Multipathing Enabled” on page 77 for specific information.

Note – If you use the `format` command when the STMS is enabled, as shown in FIGURE 4-1, you see only one instance of a device identifier for each LUN. Without the STMS, you see one identifier for each path. The `format` command is shown in FIGURE 4-1 but is not described in the manual because it is an established utility.

▼ To Add a Storage Device

1. **Create the LUN or LUNs desired.**
2. **Apply LUN masking for HBA control.**
3. **Create port-based or WWN zones on the switch.**
4. **Connect the storage device to the system.**
5. **If necessary, configure all paths to the storage device using the `cfgadm -c configure` command on all the host bus adapters that have a path to the storage device.**

The `cfgadm -c configure` command creates device nodes. This step is necessary if the storage device is connected in an NS zone and is accessed by a host port connected to a switch F port.

6. **Run the `fsck` or `newfs` commands on the device, if used for file systems.**
7. **Mount any existing file systems available on the storage device's LUNs or disk groups.**

You might need to run the `fsck` command to repair any errors in the LUNs listed in the `/etc/vfstab` file.

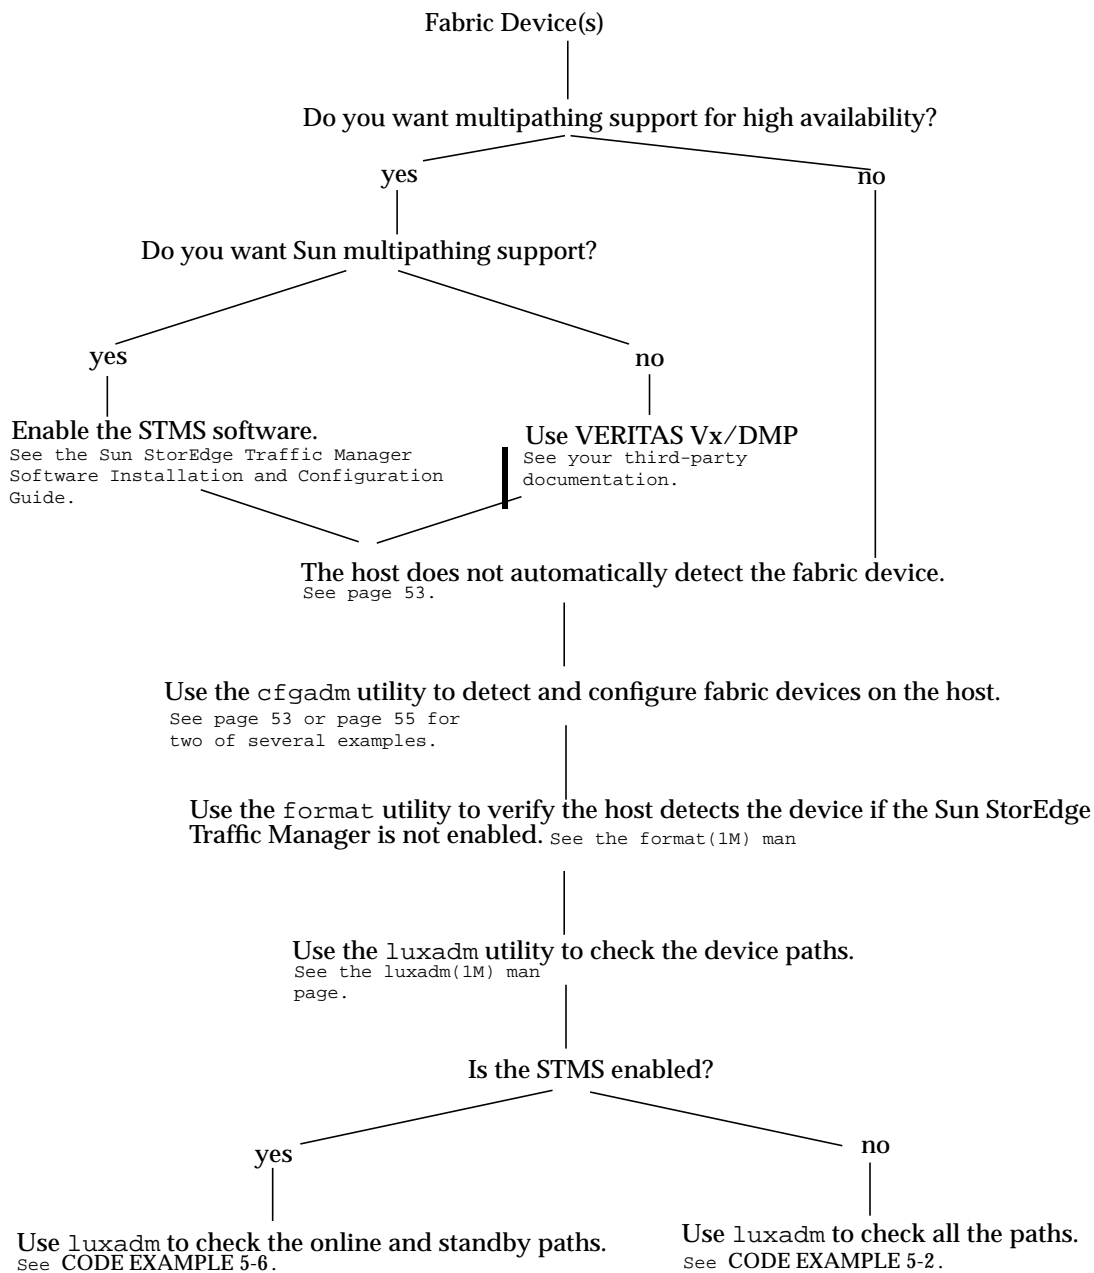


FIGURE 4-1 Decision Tree For Adding Fabric Devices

Removing a Storage Device

- If the STMS is not enabled, read the section “Creating and Removing Individual Device Nodes Without Multipathing Enabled” on page 52. See “To Unconfigure a Fabric Device” on page 59 for specific information.
- If the STMS is enabled, read the section “Creating and Removing Individual Device Nodes With Multipathing Enabled” on page 65. See “To Unconfigure a Fabric Device Associated With Multipathing Enabled Devices” on page 72 for specific information.

Consider the following:

You cannot unconfigure a single LUN for a storage target device without interfering with the access to other LUNs of the same target. In this case, you must:

1. Stop all I/O operations to the affected LUNs.
2. Release all LUNs from volume management control.
3. Unconfigure the target.
4. Disable access to the LUN to be placed offline.
5. Configure the target and re-enable access to all remaining LUNs.

Note – If you cannot disable access to the other LUNs of the same target device, disable access to this LUN (that is, instead of unconfiguring a target device with other LUNs remaining configured) and perform a reconfiguration reboot.

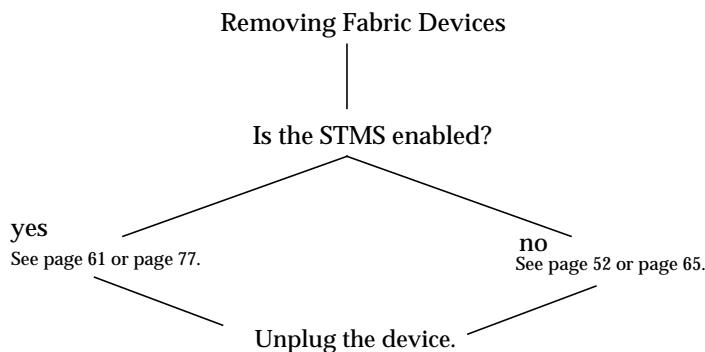


FIGURE 4-2 Decision Tree For Removing Fabric Devices

▼ To Remove a Storage Device

1. **Stop all LUN activity to the storage device.**
2. **Unmount any file systems currently using the storage device's LUNs.**
3. **If using the LUNs with any kind of volume management software, release access to these LUNs. See the documentation for your volume manager.**
4. **Unconfigure paths to the storage device using the `cfgadm -c unconfigure` command.**

This step is necessary if the storage device is connected in a nameserver zone and is accessed by a host port connected to an F port on the switch.

Fabric Device Node Recognition

After you configure the hardware in your direct-attach system or SAN, you must ensure the hosts recognize the switches and devices. This chapter explains host recognition of fabric devices, also known as 24-bit Fibre Channel addressing devices on the SAN. After configuring the devices, ports and zones in your SAN, you need to make sure that the host is aware of the devices and their switch connections. You can have up to 16 million fabric devices connected together on a SAN with Fibre Channel support.

This chapter provides generic instructions for adding and removing all supported devices in this release. This chapter includes:

- “Creating and Removing Individual Device Nodes Without Multipathing Enabled” on page 52
- “Creating and Removing Multiple Device Nodes Without Multipathing Enabled” on page 61
- “Creating and Removing Individual Device Nodes With Multipathing Enabled” on page 65
- “Creating and Removing Multiple Device Nodes With Multipathing Enabled” on page 77

There are a variety of ways in which you can ensure the host recognizes storage devices you add on your SAN. The decision trees in FIGURE 4-1 and FIGURE 4-2 help guide you through the process.

Note – You can connect a Sun StorEdge T3 array as an FC-AL device to a switch, but the Sun StorEdge host bus adapters supported in the SAN Foundation release recognize the array as a fabric device.

The scope of this chapter is limited to the operations required from the perspective of the Solaris operating environment. It does *not* cover other aspects, such as device availability and device-specific management. If devices are managed by other software, such as a volume manager, refer to the volume manager product documentation for additional instructions.

Creating and Removing Individual Device Nodes Without Multipathing Enabled

This section describes fabric device configuration tasks on a host that does not have the STMS enabled.

The procedures in this section use specific devices as examples to illustrate how to use the `cfgadm(1M)` command to detect and configure fabric devices.

The devices attached to the fabric-connected host port are not configured by default, thus those devices are not available to the host using the Solaris operating environment. Use the `cfgadm(1M)` `configure` and `unconfigure` commands to manage device node creation for fabric devices. See the `cfgadm_fp(1M)` man page for additional information.

The procedures in this section show how to detect fabric devices that are visible on a host and to configure and make them available to a host using the Solaris operating environment.

The device information that you supply and that is displayed by the `cfgadm(1M)` command depends on your system configuration.

This section contains the following topics:

- “To Detect Fabric Devices Visible on a Host” on page 53
- “Ensuring LUN Level Information Is Visible” on page 54
- “To Configure a Fabric Device Without Multipathing Enabled On the Host” on page 55
- “To Unconfigure a Fabric Device” on page 59
- “To Configure All Fabric Devices on a Fabric-Connected Host Port” on page 61
- “To Unconfigure All Fabric Devices on a Fabric-Connected Host Port” on page 63

▼ To Detect Fabric Devices Visible on a Host

This procedure uses Fibre Channel host ports `c0` and `c1`, and the devices attached to them, to provide an example of detecting fabric devices. This procedure also shows the device configuration information that is displayed with the `cfgadm(1M)` command.

Note – If you do not install the proper software patches and packages for this release, Fibre Channel devices do not display in the `cfgadm(1M)` command output. In the following examples, only failover path attachment point IDs (Ap_Ids) are listed. The Ap_Ids displayed on your system depend on your system configuration.

1. Become superuser.
2. Display the information about the attachment points on the system.

```
# cfgadm -l
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric    connected   unconfigured unknown
c1             fc-private   connected   configured  unknown
```

In this example, `c0` represents a fabric-connected host port, and `c1` represents a private, loop-connected host port. Use the `cfgadm(1M)` command to manage the device configuration on fabric-connected host ports.

By default, the device configuration on private, loop-connected host ports is managed by a host using the Solaris operating environment.

3. Display information about the host ports and their attached devices.

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric    connected   unconfigured unknown
c0::50020f2300006077 disk         connected   unconfigured unknown
c0::50020f23000063a9 disk         connected   unconfigured unknown
c0::50020f2300005f24 disk         connected   unconfigured unknown
c0::50020f2300006107 disk         connected   unconfigured unknown
c1             fc-private   connected   configured  unknown
c1::220203708b69c32b disk         connected   configured  unknown
c1::220203708ba7d832 disk         connected   configured  unknown
c1::220203708b8d45f2 disk         connected   configured  unknown
c1::220203708b9b20b2 disk         connected   configured  unknown
```

Note – The `cfgadm -l` command displays information about Fibre Channel host ports. Also use the `cfgadm -al` command to display information about Fibre Channel devices. The lines that include a port World Wide Name (WWN) in the `Ap_Id` field associated with `c0` represent a fabric device. Use the `cfgadm configure` and `unconfigure` commands to manage those devices and make them available to hosts using the Solaris operating environment. The `Ap_Id` devices with port WWNs under `c1` represent private-loop devices that are configured through the `c1` host port.

To identify which device nodes represent the same storage device on your own system, for example, log in to your Sun StorEdge T3 array and use `port list` to list the WWNs of the array controllers attached to your SAN.

Ensuring LUN Level Information Is Visible

If you issue the `cfgadm -al -o show_FCP_dev <controller_id>` command immediately after a system boots up, the output might not show the Fibre Channel Protocol (FCP) SCSI LUN level information. The information does not appear because the storage device drivers, such as the `ssd` and `st` driver, are not loaded on the running system. Use the `modinfo` command to check if the drivers are loaded. After you load the drivers, the LUN level information is visible in the `cfgadm` output. Also, issue the `format` command to confirm the disks and LUNs are visible.

▼ To Configure a Fabric Device Without Multipathing Enabled On the Host

This procedure describes how to configure a fabric device that is attached to the fabric-connected host port c0.

1. **Become superuser.**
2. **Identify the device to be configured.**

Only devices on a fabric-connected host port can be configured.

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric     connected   unconfigured unknown
c0::50020f2300006077 disk         connected   unconfigured unknown
c0::50020f23000063a9 disk         connected   unconfigured unknown
c0::50020f2300005f24 disk         connected   unconfigured unknown
c0::50020f2300006107 disk         connected   unconfigured unknown
c1             fc-private    connected   configured  unknown
c1::220203708b69c32b disk         connected   configured  unknown
c1::220203708ba7d832 disk         connected   configured  unknown
c1::220203708b8d45f2 disk         connected   configured  unknown
c1::220203708b9b20b2 disk         connected   configured  unknown
```

3. **Configure the fabric device.**

```
# cfgadm -c configure c0::50020f2300006077
```

4. **Verify that the selected fabric device is configured.**

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric     connected   configured  unknown
c0::50020f2300006077 disk         connected   configured  unknown
c0::50020f23000063a9 disk         connected   unconfigured unknown
c0::50020f2300005f24 disk         connected   unconfigured unknown
c0::50020f2300006107 disk         connected   unconfigured unknown
c1             fc-private    connected   configured  unknown
c1::220203708b69c32b disk         connected   configured  unknown
c1::220203708ba7d832 disk         connected   configured  unknown
c1::220203708b8d45f2 disk         connected   configured  unknown
c1::220203708b9b20b2 disk         connected   configured  unknown
```

Notice that the Occupant column for both `c0` and `c0::50020f2300006077` displays as configured, indicating that the `c0` port has a configured occupant and that the `c0::50020f2300006077` device is configured.

Use the `show_FCP_dev` option to display FCP SCSI LUN information for multi-LUN SCSI devices. From CODE EXAMPLE 5-1, the physical devices connected through `ap_id c0:50020f2300006077` have 4 LUNs configured.

CODE EXAMPLE 5-1 `show_FCP_dev` Output Showing Two LUNs

```
# cfgadm -al -o show_FCP_dev c0
```

Ap_Id	Type	Receptacle	Occupant	Condition
c2	fc-fabric	connected	configured	unknown
c0::50020f2300006077,0	disk	connected	configured	unknown
c0::50020f2300006077,1	disk	connected	configured	unknown
c0::50020f2300006077,2	disk	connected	configured	unknown
c0::50020f2300006077,3	disk	connected	configured	unknown

The device is now available on the host using the Solaris operating environment. CODE EXAMPLE 5-2 is an example of the `luxadm(1M)` output. Notice that four devices are listed under Paths:

```
/dev/rdsd/c0t50020f2300006077d3s2
/dev/rdsd/c0t50020f2300006077d2s2
/dev/rdsd/c0t50020f2300006077d1s2
/dev/rdsd/c0t50020f2300006077d0s2
```

The paths represent each SCSI LUN in the physical device represented by `c0::50020f2300006077`.

CODE EXAMPLE 5-2 `luxadm` Output For Four Devices and a Single Array

```
# luxadm display 50020f2300006077
```

DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port A):	O.K.
Vendor:	SUN
Product ID:	T300
WWN(Node):	50020f2000006077
WWN(Port A):	50020f2300006077
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0

CODE EXAMPLE 5-2 luxadm Output For Four Devices and a Single Array (Continued)

Device Type:	Disk device
Path(s):	/dev/rdsk/c0t50020F2300006077d3s2
/devices/pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0/ssd@w50020f2300006077,3:c,raw	
DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port A):	O.K.
Vendor:	SUN
Product ID:	T300
WWN(Node):	50020f2000006077
WWN(Port A):	50020f2300006077
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0
Device Type:	Disk device
Path(s):	/dev/rdsk/c0t50020F2300006077d2s2
/devices/pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0/ssd@w50020f2300006077,2:c,raw	
DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port A):	O.K.
Vendor:	SUN
Product ID:	T300
WWN(Node):	50020f2000006077
WWN(Port A):	50020f2300006077
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0
Device Type:	Disk device
Path(s):	/dev/rdsk/c0t50020F2300006077d1s2
/devices/pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0/ssd@w50020f2300006077,1:c,raw	
DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port B):	O.K.
Vendor:	SUN

CODE EXAMPLE 5-2 luxadm Output For Four Devices and a Single Array (Continued)

Product ID:	T300
WWN(Node):	50020f2000006077
WWN(Port B):	50020f2300006077
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0
Device Type:	Disk device
Path(s):	
	/dev/rdisk/c0t50020F2300006077d0s2
	/devices/pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0/ssd@w50020f2300006077,0:c,raw

▼ To Unconfigure a Fabric Device

This procedure describes how to unconfigure a fabric device that is attached to the fabric-connected host port `c0`.

Note – Before you unconfigure a fabric device, stop all activity to the device and unmount any file systems on the fabric device. See the administration documentation for the Solaris operating environment for unmounting instructions. If the device is under any volume manager's control, see the documentation for your volume manager for maintaining the fabric device

1. **Become superuser.**
2. **Identify the device to be unconfigured.**

Only devices on a fabric-connected host port can be unconfigured.

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric    connected   configured  unknown
c0::50020f2300006077  disk        connected   configured  unknown
c0::50020f23000063a9  disk        connected   configured  unknown
c0::50020f2300005f24  disk        connected   configured  unknown
c0::50020f2300006107  disk        connected   configured  unknown
c1             fc-private   connected   configured  unknown
c1::220203708b69c32b  disk        connected   configured  unknown
c1::220203708ba7d832  disk        connected   configured  unknown
c1::220203708b8d45f2  disk        connected   configured  unknown
c1::220203708b9b20b2  disk        connected   configured  unknown
```

3. **Unconfigure the fabric device.**

```
# cfgadm -c unconfigure c0::50020f2300006077
```

4. Verify that the selected fabric device is unconfigured.

```
# cfgadm -al
Ap_Id                Type                Receptacle  Occupant  Condition
c0                   fc-fabric          connected   configured unknown
c0::50020f2300006077 disk                connected   unconfigured unknown
c0::50020f23000063a9 disk                connected   configured  unknown
c0::50020f2300005f24 disk                connected   configured  unknown
c0::50020f2300006107 disk                connected   configured  unknown
c1                   fc-private         connected   configured  unknown
c1::220203708b69c32b disk                connected   configured  unknown
c1::220203708ba7d832 disk                connected   configured  unknown
c1::220203708b8d45f2 disk                connected   configured  unknown
c1::220203708b9b20b2 disk                connected   configured  unknown
```

Creating and Removing Multiple Device Nodes Without Multipathing Enabled

Procedures for creating and removing multiple devices are similar to those described in “Creating and Removing Individual Device Nodes Without Multipathing Enabled” on page 52. This section explains the finer differences. Make sure you first identify the devices visible to the host with the procedure “To Detect Fabric Devices Visible on a Host” on page 53. This section covers:

- “To Configure All Fabric Devices on a Fabric-Connected Host Port” on page 61
- “To Unconfigure All Fabric Devices on a Fabric-Connected Host Port” on page 63

▼ To Configure All Fabric Devices on a Fabric-Connected Host Port

This procedure describes how to configure all unconfigured fabric devices that are attached to a fabric-connected host port. The port used as an example is `c0`.

1. **Become superuser.**
2. **Identify the devices to be configured.**

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric    connected   unconfigured unknown
c0::50020f2300006077 disk         connected   unconfigured unknown
c0::50020f23000063a9 disk         connected   unconfigured unknown
c0::50020f2300005f24 disk         connected   unconfigured unknown
c0::50020f2300006107 disk         connected   unconfigured unknown
c1             fc-private   connected   configured   unknown
c1::220203708b69c32b disk         connected   configured   unknown
c1::220203708ba7d832 disk         connected   configured   unknown
c1::220203708b8d45f2 disk         connected   configured   unknown
c1::220203708b9b20b2 disk         connected   configured   unknown
```

3. **Configure all of the unconfigured devices on the selected port.**

```
# cfgadm -c configure c0
```

Note – This operation repeats the `configure` operation of an individual device for all the devices on `c0`, and can be time consuming if the number of devices on `c0` is large.

4. Verify that all devices on `c0` are configured.

```
# cfgadm -al
Ap_Id                Type                Receptacle  Occupant  Condition
c0                   fc-fabric          connected   configured unknown
c0::50020f2300006077 disk                connected   configured unknown
c0::50020f23000063a9 disk                connected   configured unknown
c0::50020f2300005f24 disk                connected   configured unknown
c0::50020f2300006107 disk                connected   configured unknown
c1                   fc-private         connected   configured unknown
c1::220203708b69c32b disk                connected   configured unknown
c1::220203708ba7d832 disk                connected   configured unknown
c1::220203708b8d45f2 disk                connected   configured unknown
c1::220203708b9b20b2 disk                connected   configured unknown
```

The `show_FCP_dev` option displays FCP SCSI LUN information for multiple LUN SCSI devices. From the example in CODE EXAMPLE 5-3, the physical devices represented by `c0::50020f2300006077` and `c0::50020f2300006107` have four LUNs configured respectively. The physical devices represented by `c0::50020f23000063a9` and `c0::50020f2300005f24` have two LUNs configured respectively.

CODE EXAMPLE 5-3 `show_FCP_dev` Output For Multiple LUNs and Two Devices

```
# cfgadm -al -o show_FCP_dev c0
Ap_Id                Type                Receptacle  Occupant  Condition
c0                   fc-fabric          connected   configured unknown
c0::50020f2300006077,0 disk                connected   configured unknown
c0::50020f2300006077,1 disk                connected   configured unknown
c0::50020f2300006077,2 disk                connected   configured unknown
c0::50020f2300006077,3 disk                connected   configured unknown
c0::50020f23000063a9,0 disk                connected   configured unknown
c0::50020f23000063a9,1 disk                connected   configured unknown
c0::50020f2300005f24,0 disk                connected   configured unknown
c0::50020f2300005f24,1 disk                connected   configured unknown
c0::50020f2300006107,0 disk                connected   configured unknown
c0::50020f2300006107,1 disk                connected   configured unknown
c0::50020f2300006107,2 disk                connected   configured unknown
c0::50020f2300006107,3 disk                connected   configured unknown
```

▼ To Unconfigure All Fabric Devices on a Fabric-Connected Host Port

This procedure describes how to unconfigure all configured fabric devices that are attached to a fabric-connected host port.

1. **Become superuser.**
2. **Identify the fabric devices to be unconfigured.**

Only devices on a fabric-connected host port can be unconfigured.

```
# cfgadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	fc-fabric	connected	configured	unknown
c0::50020f2300006077	disk	connected	configured	unknown
c0::50020f23000063a9	disk	connected	configured	unknown
c0::50020f2300005f24	disk	connected	configured	unknown
c0::50020f2300006107	disk	connected	configured	unknown
c1	fc-private	connected	configured	unknown
c1::220203708b69c32b	disk	connected	configured	unknown
c1::220203708ba7d832	disk	connected	configured	unknown
c1::220203708b8d45f2	disk	connected	configured	unknown
c1::220203708b9b20b2	disk	connected	configured	unknown

3. **Unconfigure all of the configured fabric devices on a selected port.**

Note – Stop all activity to each fabric device on the selected port and unmount any file systems on each fabric device. If the device is under any volume manager’s control, see the documentation for your volume manager for maintaining the fabric device.

```
# cfgadm -c unconfigure c0
```

Note – This operation repeats the `unconfigure` operation of an individual device for all the devices on `c0` and it can be time-consuming if the number of devices on `c0` is large.

4. Verify that all the devices on c0 are unconfigured.

```
# cfgadm -al
Ap_Id                Type                Receptacle  Occupant  Condition
c0                   fc-fabric          connected   unconfigured unknown
c0::50020f2300006077 disk                connected   unconfigured unknown
c0::50020f23000063a9 disk                connected   unconfigured unknown
c0::50020f2300005f24 disk                connected   unconfigured unknown
c0::50020f2300006107 disk                connected   unconfigured unknown
c1                   fc-private         connected   configured  unknown
c1::220203708b69c32b disk                connected   configured  unknown
c1::220203708ba7d832 disk                connected   configured  unknown
c1::220203708b8d45f2 disk                connected   configured  unknown
c1::220203708b9b20b2 disk                connected   configured  unknown
```

Notice that the Occupant column of c0 and all the fabric devices attached to it are displayed as unconfigured.

Creating and Removing Individual Device Nodes With Multipathing Enabled

This section describes how to perform fabric device configuration steps on a host that has the Sun StorEdge Traffic Manager multipathing software enabled.

The devices that are attached to fabric-connected HBA ports are not configured by default. These devices are thus not available to the host using the Solaris operating environment when a host port is initially connected to a fabric. The procedures in this section illustrate steps to detect fabric devices that are visible on a host and to configure them as Sun StorEdge Traffic Manager devices to make them available to the host using the Solaris operating environment.

The device information that you supply, and that which is displayed by the `cfgadm(1M)` command, depends on your system configuration. (For more information on the `cfgadm` command, see the `cfgadm_fp(1M)` and `cfgadm(1M)` man pages.)

This section contains the following procedures:

- “To Detect Fabric Devices Visible to a Host” on page 66
- “Ensuring LUN Level Information Is Visible” on page 67
- “To Configure Fabric-Connected Device Nodes With Multipathing Enabled Devices” on page 68
- “To Unconfigure a Fabric Device Associated With Multipathing Enabled Devices” on page 72
- “To Unconfigure One Path to a Multipathed Device” on page 74
- “To Configure All Fabric-Connected Devices With Multipathing Enabled” on page 78
- “To Unconfigure All Fabric-Connected Devices With Multipathing Enabled” on page 83

▼ To Detect Fabric Devices Visible to a Host

This procedure shows Fibre Channel host ports `c0`, `c1`, and `c2` and the devices attached to them. It illustrates fabric device detection and device configuration using the `cfgadm(1M)` command.

Note – If the proper `cfgadm` support for Fibre Channel devices is not installed, Fibre Channel devices do not display in the `cfgadm(1M)` command output. In the following examples, only failover path attachment points (`Ap_Ids`) are listed. The `Ap_Ids` displayed on your system depend on your system configuration.

1. Become superuser.
2. Display information about the attachment points on the system.

```
# cfgadm -l
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric     connected   unconfigured unknown
c1             fc-private    connected   configured   unknown
c2             fc-fabric     connected   unconfigured unknown
```

An `Ap_Id` on a fabric-connected host port is a path to a Sun StorEdge Traffic Manager device. In this example, `c0` and `c2` represent fabric-connected host ports. Also, `c1` represents a private, loop-connected host port. Use the `cfgadm(1M)` command to manage the device configuration on fabric-connected host ports.

By default, the device configuration on private, loop-connected host ports are managed by a host using the Solaris operating environment.

3. Display information about the host ports and their attached devices.

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric     connected   unconfigured unknown
c0::50020f2300006077 disk          connected   unconfigured unknown
c0::50020f23000063a9 disk          connected   unconfigured unknown
c1             fc-private    connected   configured   unknown
c1::220203708b69c32b disk          connected   configured   unknown
c1::220203708ba7d832 disk          connected   configured   unknown
c1::220203708b8d45f2 disk          connected   configured   unknown
c1::220203708b9b20b2 disk          connected   configured   unknown
c2             fc-fabric     connected   unconfigured unknown
c2::50020f2300005f24 disk          connected   unconfigured unknown
c2::50020f2300006107 disk          connected   unconfigured unknown
```

Note – The `cfgadm -l` command displays information about Fibre Channel host ports. Also use the `cfgadm -al` command to display information about Fibre Channel devices. The lines that include a port World Wide Name (WWN) in the `Ap_Id` field associated with `c0` represent a fabric device. Use the `cfgadm configure` and `unconfigure` commands to manage those devices and make available to hosts using the Solaris operating environment. The `Ap_Id` devices with port WWNs under `c1` represent private-loop devices that are configured through the `c1` host port.

In the previous example, host ports `c0` and `c2` are connected to the two ports of a Sun StorEdge T3 enterprise array that has two port WWNs associated with it: `Ap_Ids 50020f2300006077` and `50020f2300006107`.

If your storage device is a Sun StorEdge T3 array and you want to identify which device nodes represent the same storage device on your own system, log in to the array and use `port list` to list the WWNs of the array controllers attached to your SAN.

Ensuring LUN Level Information Is Visible

If you issue the `cfgadm -al -o show_FCP_dev <controller_id>` command immediately after a system boots up, the output might not show the Fibre Channel Protocol (FCP) SCSI LUN level information. The information does not appear because the storage device drivers, such as the `ssd` and `st` driver, are not loaded on the running system. Use the `modinfo` command to check if the drivers are loaded. After you load the drivers, the LUN level information is visible in the `cfgadm` output.

▼ To Configure Fabric-Connected Device Nodes With Multipathing Enabled Devices

This procedure uses fabric-connected host ports `c0` and `c2` to configure fabric devices as Sun StorEdge Traffic Manager devices on a host that has the STMS enabled.

Note – Whether the STMS is enabled or not, the `cfgadm -c configure` command for fabric devices is identical but the result is different. When the Sun Storage Traffic Manager software is enabled, the host using the Solaris operating environment creates device node and path information that includes STMS information. The devices are still listed as `fc-fabric` devices, but the specific `Ap_Ids` are marked as configured in the Occupant column.

1. Become superuser.
2. Identify the port WWN of the device to be configured as a Sun StorEdge Traffic Manager device.

Look for devices on a fabric-connected host port, marked as `fc-fabric`. These are the devices you can configure with the `cfgadm -c configure` command.

CODE EXAMPLE 5-4 `cfgadm` Listing of fabric and Private-Loop Devices

```
# cfgadm -al
Ap_Id                Type           Receptacle      Occupant        Condition
c0                   fc-fabric      connected        unconfigured    unknown
c0::50020f2300006077 disk           connected        unconfigured    unknown
c0::50020f23000063a9 disk           connected        unconfigured    unknown
c1                   fc-private     connected        configured       unknown
c1::220203708b69c32b disk           connected        configured       unknown
c1::220203708ba7d832 disk           connected        configured       unknown
c1::220203708b8d45f2 disk           connected        configured       unknown
c1::220203708b9b20b2 disk           connected        configured       unknown
c2                   fc-fabric      connected        unconfigured    unknown
c2::50020f23000005f24 disk           connected        unconfigured    unknown
c2::50020f23000006107 disk           connected        unconfigured    unknown
```

In CODE EXAMPLE 5-4, the `c0::50020f2300006077` and `c2::50020f2300006107` `Ap_Ids` represent the same storage device with different port WWNs for the storage device controllers. The `c0` and `c2` host ports are enabled for use by the STMS.

3. Configure the fabric device and make Sun StorEdge Traffic Manager devices available to the host.

```
# cfgadm -c configure c0::50020f2300006077 c2::50020f2300006107
```

4. Verify that the selected devices are configured.

```
# cfgadm -al
Ap_Id                Type           Receptacle    Occupant      Condition
c0                   fc-fabric     connected     configured    unknown
c0::50020f2300006077 disk          connected     configured    unknown
c0::50020f23000063a9 disk          connected     unconfigured  unknown
c1                   fc-private    connected     configured    unknown
c1::220203708b69c32b disk          connected     configured    unknown
c1::220203708ba7d832 disk          connected     configured    unknown
c1::220203708b8d45f2 disk          connected     configured    unknown
c1::220203708b9b20b2 disk          connected     configured    unknown
c2                   fc-fabric     connected     configured    unknown
c2::50020f2300005f24 disk          connected     unconfigured  unknown
c2::50020f2300006107 disk          connected     configured    unknown
```

Notice that the Occupant column of `c0` and `c0::50020f2300006077` specifies `configured`, which indicates that the `c0` port has at least one configured occupant and that the `c0::50020f2300006077` device is configured. The same change has been made in `c2` and `c2::50020f2300006107`.

After completing the configure operation without an error, STMS enabled devices are created on the host using the Solaris operating environment. If the physical device represented by `c0::50020f2300006077` and `c2::50020f2300006107` has multiple SCSI LUNs configured, each LUN is configured as a Sun StorEdge Traffic Manager device. CODE EXAMPLE 5-5 shows that two LUNs are configured through `c0::50020f2300006077` and `c2::50020f2300006107`. Each `ap_id` is associated with a path to those Sun StorEdge Traffic Manager devices.

CODE EXAMPLE 5-5 show_FCP_dev Output For Two LUNs On a Device

```
# cfgadm -al -o show_FCP_dev c0::50020f2300006077
c2::50020f2300006107
Ap_Id                Type           Receptacle    Occupant      Condition
c0::50020f2300006077,0 disk          connected     configured    unknown
c0::50020f2300006077,1 disk          connected     configured    unknown
c2::50020f2300006107,0 disk          connected     configured    unknown
c2::50020f2300006107,1 disk          connected     configured    unknown
```

In CODE EXAMPLE 5-6, notice that two STMS enabled devices

```
/dev/rdsk/c6t60020F20000061073AC8B52D000B74A3d0s2
/dev/rdsk/c6t60020F20000061073AC8B4C50004ED3Ad0s2
```

are created for the device represented by c0::50020f2300006077 and c2::50020f2300006107.

CODE EXAMPLE 5-6 luxadm(1M) Output For Multiple Device Nodes With the STMS Enabled

# luxadm display 50020f2300006077	
DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port A):	O.K.
Status(Port B):	O.K.
Vendor:	SUN
Product ID:	T300
WWN(Node):	50020f2000006077
WWN(Port A):	50020f2300006077
WWN(Port B):	50020f2300006107
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0
Device Type:	Disk device
Path(s):	
/dev/rdsk/c6t60020F20000061073AC8B52D000B74A3d0s2	
/devices/scsi_vhci/ssd@g60020f20000061073ac8b52d000b74a3:c,raw	
Controller	/devices/pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0
Device Address	50020f2300006107,1
Class	secondary
State	STANDBY
Controller	/devices/pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,0
Device Address	50020f2300006077,1
Class	primary
State	ONLINE
DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port A):	O.K.
Status(Port B):	O.K.
Vendor:	SUN
Product ID:	T300
WWN(Node):	50020f2000006107

CODE EXAMPLE 5-6 luxadm(1M) Output For Multiple Device Nodes With the STMS Enabled (Continued)

WWN(Port A):	50020f2300006107
WWN(Port B):	50020f2300006077
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0
Device Type:	Disk device
Path(s):	
/dev/rdisk/c6t60020F20000061073AC8B4C50004ED3Ad0s2	
/devices/scsi_vhci/ssd@g60020f20000061073ac8b4c50004ed3a:c,raw	
Controller	/devices/pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0
Device Address	50020f2300006107,0
Class	primary
State	ONLINE
Controller	/devices/pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,0
Device Address	50020f2300006077,0
Class	secondary
State	STANDBY

Note – The luxadm (1M) output on device 50020f2300006107 shows the same information as the previous display.

▼ To Unconfigure a Fabric Device Associated With Multipathing Enabled Devices

This procedure shows fabric-connected host ports `c0` and `c2` to illustrate how to unconfigure fabric devices associated with Sun StorEdge Traffic Manager devices.

Note – Whether the STMS is enabled or not, the `cfgadm -c unconfigure` command for Fabric devices is identical, but the result is different. When the Sun Storage Traffic Manager software is enabled, the host using the Solaris operating environment creates device-node and path information that includes STMS information. The devices are still listed as `fc-fabric` devices, but the specific `Ap_Ids` are marked as `unconfigured` in the Occupant column.

1. Become superuser.
2. Identify the port WWN of the fabric device to be unconfigured.

```
# cfgadm -al
Ap_Id                Type           Receptacle  Occupant  Condition
c0                   fc-fabric      connected   configured unknown
c0::50020f2300006077 disk           connected   configured unknown
c0::50020f23000063a9 disk           connected   configured unknown
c1                   fc-private     connected   configured unknown
c1::220203708b69c32b disk           connected   configured unknown
c1::220203708ba7d832 disk           connected   configured unknown
c1::220203708b8d45f2 disk           connected   configured unknown
c1::220203708b9b20b2 disk           connected   configured unknown
c2                   fc-fabric      connected   configured unknown
c2::50020f23000005f24 disk           connected   configured unknown
c2::50020f2300006107 disk           connected   configured unknown
```

In this example, the `c0::50020f2300006077` and `c2::50020f2300006107` `Ap_Ids` represent different port WWNs for the same device associated with a Sun StorEdge Traffic Manager device. The `c0` and `c2` host ports are enabled for use by the STMS.

Note – Stop all device activity to each fabric device on the selected port and unmount any file systems on each fabric device. If the device is under any volume manager's control, see the documentation for your volume manager for maintaining the fabric device.

3. Unconfigure fabric devices associated with the Sun StorEdge Traffic Manager device.

Only devices on a fabric-connected host port can be unconfigured through the `cfgadm -c unconfigure` command.

```
# cfgadm -c unconfigure c0::50020f2300006077 c2::50020f2300006107
```

Note – You can remove a device from up to eight paths individually, as in the example command `cfgadm -c unconfigure c0::1111, c1::2222, c3::3333`, etc. As an alternative, you can remove an entire set of paths from the host, as in the example `cfgadm -c unconfigure c0`.

4. Verify that the selected devices are unconfigured.

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric    connected   configured  unknown
c0::50020f2300006077 disk         connected   unconfigured unknown
c0::50020f23000063a9 disk         connected   configured  unknown
c1             fc-private   connected   configured  unknown
c1::220203708b69c32b disk         connected   configured  unknown
c1::220203708ba7d832 disk         connected   configured  unknown
c1::220203708b8d45f2 disk         connected   configured  unknown
c1::220203708b9b20b2 disk         connected   configured  unknown
c2             fc-fabric    connected   configured  unknown
c2::50020f2300005f24 disk         connected   configured  unknown
c2::50020f2300006107 disk         connected   unconfigured unknown
```

Notice that the Ap_Ids `c0::50020f2300006077` and `c2::50020f2300006107` are unconfigured. The Occupant column of `c0` and `c2` still displays those ports as configured because they have other configured occupants.

The Sun StorEdge Traffic Manager devices associated with the Ap_Ids `c0::50020f2300006077` and `c2::50020f2300006107` are no longer available to the host using the Solaris operating environment. The two Sun StorEdge Traffic Manager devices,

```
/dev/rdsk/c6t60020F20000061073AC8B52D000B74A3d0s2
```

and

```
/dev/rdsk/c6t60020F20000061073AC8B4C50004ED3Ad0s2
```

are removed from the host.

▼ To Unconfigure One Path to a Multipathed Device

In “To Unconfigure a Fabric Device Associated With Multipathing Enabled Devices” on page 72, a storage device is connected to the host using the Solaris operating environment by two Ap_Ids, `c0::50020f2300006077` and `c2::50020f2300006107`. For the STMS enabled host, each Ap_Id is associated with the path to a Sun StorEdge Traffic Manager device that represents the physical storage device. This procedure shows how to unconfigure a device associated with `c2::50020f2300006107` and leave the other Ap_Id, `50020f2300006077`, configured.

1. **Become superuser.**
2. **Identify the Ap_Id of the Sun StorEdge Traffic Manager device to be unconfigured.**

An Ap_Id on a fabric-connected host port is a path to a Sun StorEdge Traffic Manager device. Only devices on a fabric-connected host port can be unconfigured through the `cfgadm unconfigure` command.

CODE EXAMPLE 5-7 `cfgadm` Listing of Port WWNs For Physical Devices To Unconfigure

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant  Condition
c0             fc-fabric     connected   configured unknown
c0::50020f2300006077 disk         connected   configured unknown
c0::50020f23000063a9 disk         connected   configured unknown
c1             fc-private    connected   configured unknown
c1::220203708b69c32b disk         connected   configured unknown
c1::220203708ba7d832 disk         connected   configured unknown
c1::220203708b8d45f2 disk         connected   configured unknown
c1::220203708b9b20b2 disk         connected   configured unknown
c2             fc-fabric     connected   configured unknown
c2::50020f2300005f24 disk         connected   configured unknown
c2::50020f2300006107 disk         connected   configured unknown
```

In CODE EXAMPLE 5-7, `c0::50020f2300006077` and `c2::50020f2300006107` Ap_Ids represent different port WWNs for the same device.

3. **Unconfigure the Ap_Id associated with Sun StorEdge Traffic Manager device.**

Note – If the Ap_Id represents the last configured path to the Sun StorEdge Traffic Manager device, stop all activity to the path and unmount any file systems on it. If the multipathed device is under any volume manager’s control, see the documentation for your volume manager for maintaining the fabric device

In the example that follows, the path represented as `c2::50020f2300006107` is unconfigured, and `c0::50020f2300006077` remains configured to show how you can unconfigure just one of multiple paths for a multipathed device.

```
# cfgadm -c unconfigure c2::50020f2300006107
```

4. Verify that the selected path `c2::50020f2300006107` is unconfigured.

```
# cfgadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	fc-fabric	connected	configured	unknown
c0::50020f2300006077	disk	connected	configured	unknown
c0::50020f23000063a9	disk	connected	configured	unknown
c1	fc-private	connected	configured	unknown
c1::220203708b69c32b	disk	connected	configured	unknown
c1::220203708ba7d832	disk	connected	configured	unknown
c1::220203708b8d45f2	disk	connected	configured	unknown
c1::220203708b9b20b2	disk	connected	configured	unknown
c2	fc-fabric	connected	configured	unknown
c2::50020f2300005f24	disk	connected	configured	unknown
c2::50020f2300006107	disk	connected	unconfigured	unknown

The Sun StorEdge Traffic Manager devices associated with that Ap_Id are still available to a host using the Solaris operating environment through the other path, represented by `c0::50020f2300006077`. A device can be connected to multiple Ap_Id's and an Ap_Id can be connected to multiple devices.

CODE EXAMPLE 5-8 shows example output you would see for the Sun StorEdge Traffic Manager devices if you use the `luxadm(1M)` command after performing this procedure. Although the path represented by `c2::50020f2300006107` is no longer listed, the path represented by `c0::50020f2300006077` is displayed for Sun StorEdge Traffic Manager devices

```
/dev/rdsk/c6t60020F20000061073AC8B52D000B74A3d0s2
```

and

```
/dev/rdsk/c6t60020F20000061073AC8B4C50004ED3Ad0s2.
```

CODE EXAMPLE 5-8 luxadm (1M) Output For Two Device Nodes With a Single Path Available With the Sun StorEdge Traffic Manager Software Enabled

# luxadm display 50020f2300006077	
DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port A):	O.K.
Vendor:	SUN
Product ID:	T300
WWN(Node):	50020f2000006077
WWN(Port A):	50020f2300006077
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0
Device Type:	Disk device
Path(s):	
/dev/rdsk/c6t60020F20000061073AC8B52D000B74A3d0s2	
/devices/scsi_vhci/ssd@g60020f20000061073ac8b52d000b74a3:c,raw	
Controller	/devices/pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,0
Device Address	50020f2300006077,1
Class	primary
State	ONLINE
DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port B):	O.K.
Vendor:	SUN
Product ID:	T300
WWN(Node):	50020f2000006077
WWN(Port B):	50020f2300006077
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0
Device Type:	Disk device
Path(s):	
/dev/rdsk/c6t60020F20000061073AC8B4C50004ED3Ad0s2	
/devices/scsi_vhci/ssd@g60020f20000061073ac8b4c50004ed3a:c,raw	

CODE EXAMPLE 5-8 luxadm (1M) Output For Two Device Nodes With a Single Path Available With the Sun StorEdge Traffic Manager Software Enabled

Controller	/devices/pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,0
Device Address	50020f2300006077,0
Class	secondary
State	ONLINE

Creating and Removing Multiple Device Nodes With Multipathing Enabled

The procedures for creating and removing multiple device nodes is similar to those in the section “Creating and Removing Individual Device Nodes With Multipathing Enabled” on page 65. This section explains the finer differences. Make sure you have first identified the fabric devices by using the procedure “To Detect Fabric Devices Visible to a Host” on page 66 before you configure or remove device nodes. This section covers:

- “To Configure All Fabric-Connected Devices With Multipathing Enabled” on page 78
- “To Unconfigure All Fabric-Connected Devices With Multipathing Enabled” on page 83

▼ To Configure All Fabric-Connected Devices With Multipathing Enabled

In this example, an Ap_Id on a fabric-connected host port is a path to a Sun StorEdge Traffic Manager device. For example, all devices with a path through c2 are to be configured, but none through c0 are to be configured. c2 is an attachment point from the host to the fabric, whereas c2::50020f2300006107 is an attachment point from the storage to the fabric. A host detects all the storage devices in a fabric for which it is configured.

Configuring an Ap_Id on a Sun StorEdge Traffic Manager device that has already been configured through another Ap_Id results in an additional path to the previously configured device. Note that a new Solaris device is not created in this case. A Solaris device is created only the first time an Ap_Id to a corresponding Sun StorEdge Traffic Manager device is configured.

1. Become superuser.
2. Identify the fabric-connected host port to be configured.

```
# cfgadm -al
Ap_Id           Type           Receptacle    Occupant      Condition
c0              fc-fabric     connected     configured    unknown
c0::50020f2300006077 disk          connected     configured    unknown
c0::50020f23000063a9 disk          connected     configured    unknown
c1              fc-private    connected     configured    unknown
c1::220203708b69c32b disk          connected     configured    unknown
c1::220203708ba7d832 disk          connected     configured    unknown
c1::220203708b8d45f2 disk          connected     configured    unknown
c1::220203708b9b20b2 disk          connected     configured    unknown
c2              fc-fabric     connected     unconfigured  unknown
c2::50020f2300005f24 disk          connected     unconfigured  unknown
c2::50020f2300006107 disk          connected     unconfigured  unknown
```

Devices represented by Ap_Ids c0::50020f2300006077 and c2::50020f2300006107 are two paths to the same physical device, with c0::50020f2300006077 already configured. The luxadm output in CODE EXAMPLE 5-9 shows the Solaris device associated with this path. One path is configured.

CODE EXAMPLE 5-9 luxadm (1M) Output for Multiple Device Nodes Associated With One Device Path and STMS Enabled

# luxadm display 50020f2300006077	
DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port A):	O.K.
Vendor:	SUN
Product ID:	T300
WWN(Node):	50020f2000006077
WWN(Port A):	50020f2300006077
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0
Device Type:	Disk device
Path(s):	
	/dev/rdisk/c6t60020F20000061073AC8B52D000B74A3d0s2
	/devices/scsi_vhci/ssd@g60020f20000061073ac8b52d000b74a3:c,raw
Controller	/devices/pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,0
Device Address	50020f2300006077,1
Class	primary
State	ONLINE
DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port B):	O.K.
Vendor:	SUN
Product ID:	T300
WWN(Node):	50020f2000006077
WWN(Port B):	50020f2300006077
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0
Device Type:	Disk device
Path(s):	
	/dev/rdisk/c6t60020F20000061073AC8B4C50004ED3Ad0s2
	/devices/scsi_vhci/ssd@g60020f20000061073ac8b4c50004ed3a:c,raw

CODE EXAMPLE 5-9 luxadm (1M) Output for Multiple Device Nodes Associated With One Device Path and STMS Enabled (Continued)

```
Controller          /devices/pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,0
Device Address      50020f2300006077,0
Class                secondary
State                ONLINE
```

3. Configure the unconfigured devices on the selected port.

```
# cfgadm -c configure c2
```

Note – This operation repeats the `configure` command of an individual device for all the devices on `c2` and it can be time-consuming if the number of devices on `c2` is large.

4. Verify that all devices on `c2` are configured.

```
# cfgadm -al
Ap_Id           Type           Receptacle    Occupant      Condition
c0              fc-fabric     connected     configured    unknown
c0::50020f2300006077 disk          connected     configured    unknown
c0::50020f23000063a9 disk          connected     configured    unknown
c1              fc-private    connected     configured    unknown
c1::220203708b69c32b disk          connected     configured    unknown
c1::220203708ba7d832 disk          connected     configured    unknown
c1::220203708b8d45f2 disk          connected     configured    unknown
c1::220203708b9b20b2 disk          connected     configured    unknown
c2              fc-fabric     connected     configured    unknown
c2::50020f2300005f24 disk          connected     configured    unknown
c2::50020f2300006107 disk          connected     configured    unknown
```

Notice that the Occupant column of `c2` and all of the devices under `c2` is marked as configured.

The `show_FCP_dev` option displays FCP SCSI LUN information for multiple LUN SCSI devices. In CODE EXAMPLE 5-10, the physical devices connected through by `c2::50020f2300006107` and `c2::50020f2300005f24` have two LUNs configured respectively.

CODE EXAMPLE 5-10 `show_FCP_dev` Output For Multiple LUNs and Multiple Physical Devices

Ap_Id	Type	Receptacle	Occupant	Condition
c2	fc-fabric	connected	configured	unknown
c2::50020f2300005f24,0	disk	connected	configured	unknown
c2::50020f2300005f24,1	disk	connected	configured	unknown
c2::50020f2300006107,0	disk	connected	configured	unknown
c2::50020f2300006107,1	disk	connected	configured	unknown

In CODE EXAMPLE 5-11, notice that two STMS enabled devices

```
/dev/rdisk/c6t60020F20000061073AC8B52D000B74A3d0s2
/dev/rdisk/c6t60020F20000061073AC8B4C50004ED3Ad0s2
```

are created for the device represented by `c0::50020f2300006077` and `c2::50020f2300006107`.

CODE EXAMPLE 5-11 `luxadm(1M)` Output For Multiple Device Nodes With the STMS Enabled

DEVICE PROPERTIES for disk: 50020f2300006077	
Status(Port A):	O.K.
Status(Port B):	O.K.
Vendor:	SUN
Product ID:	T300
WWN(Node):	50020f2000006077
WWN(Port A):	50020f2300006077
WWN(Port B):	50020f2300006107
Revision:	0117
Serial Num:	Unsupported
Unformatted capacity:	558448.000 MBytes
Write Cache:	Enabled
Read Cache:	Enabled
Minimum prefetch:	0x0
Maximum prefetch:	0x0
Device Type:	Disk device
Path(s):	
/dev/rdisk/c6t60020F20000061073AC8B52D000B74A3d0s2	

CODE EXAMPLE 5-11 luxadm(1M) Output For Multiple Device Nodes With the STMS Enabled (Continued)

```

/devices/scsi_vhci/ssd@g60020f20000061073ac8b52d000b74a3:c,raw
Controller      /devices/pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0
  Device Address 50020f2300006107,1
  Class         secondary
  State         STANDBY
Controller      /devices/pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,0
  Device Address 50020f2300006077,1
  Class         primary
  State         ONLINE

DEVICE PROPERTIES for disk: 50020f2300006077
Status(Port A): O.K.
Status(Port B): O.K.
Vendor:         SUN
Product ID:     T300
WWN(Node):      50020f2000006107
WWN(Port A):    50020f2300006107
WWN(Port B):    50020f2300006077
Revision:       0117
Serial Num:     Unsupported
Unformatted capacity: 558448.000 MBytes
Write Cache:    Enabled
Read Cache:     Enabled
  Minimum prefetch: 0x0
  Maximum prefetch: 0x0
Device Type:    Disk device
Path(s):

/dev/rdisk/c6t60020F20000061073AC8B4C50004ED3Ad0s2
/devices/scsi_vhci/ssd@g60020f20000061073ac8b4c50004ed3a:c,raw
Controller      /devices/pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0
  Device Address 50020f2300006107,0
  Class         primary
  State         ONLINE
Controller      /devices/pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,0
  Device Address 50020f2300006077,0
  Class         secondary
  State         STANDBY

```

Note – The luxadm (1M) output on device 50020f2300006107 shows the same information as the previous display.

▼ To Unconfigure All Fabric-Connected Devices With Multipathing Enabled

Note – Whether or not the Sun StorEdge Traffic Manager multipathing software is enabled, the `cfgadm -c unconfigure` command for fabric devices is identical, but the result of the operation is different. When the STMS is enabled, the host using the Solaris operating environment removes Sun StorEdge Traffic Manager-related path or device-node information.

An `Ap_Id` on a fabric-connected host port is a path to a Sun StorEdge Traffic Manager device.

When a Sun StorEdge Traffic Manager device has multiple `Ap_Ids` connected to it, the device is still available to the host using the Solaris operating environment after you unconfigure an `Ap_Id`. After you unconfigure the last `Ap_Id`, no additional paths remain and the Sun StorEdge Manager device is unavailable to the host using the Solaris operating environment.

1. Become superuser.
2. Identify the devices to be unconfigured.

Only devices on a fabric-connected host port can be unconfigured.

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric    connected   configured  unknown
c0::50020f2300006077  disk        connected   configured  unknown
c0::50020f23000063a9  disk        connected   configured  unknown
c1             fc-private   connected   configured  unknown
c1::220203708b69c32b  disk        connected   configured  unknown
c1::220203708ba7d832  disk        connected   configured  unknown
c1::220203708b8d45f2  disk        connected   configured  unknown
c1::220203708b9b20b2  disk        connected   configured  unknown
c2             fc-fabric    connected   configured  unknown
c2::50020f2300005f24  disk        connected   configured  unknown
c2::50020f2300006107  disk        connected   configured  unknown
```

3. Unconfigure all of the configured devices on the selected port.

```
# cfgadm -c unconfigure c2
```

Note – This operation repeats the `unconfigure` command of an individual device for all devices on `c2` and it can be time-consuming if the number of devices on `c2` is large.

4. Verify that all devices on `c2` are unconfigured.

Notice that the Occupant column lists `c2` and all the devices attached to `c2` as unconfigured.

```
# cfgadm -al
Ap_Id                Type                Receptacle  Occupant  Condition
c0                   fc-fabric          connected   configured unknown
c0::50020f2300006077 disk                connected   configured unknown
c0::50020f23000063a9 disk                connected   configured unknown
c1                   fc-private         connected   configured unknown
c1::220203708b69c32b disk                connected   configured unknown
c1::220203708ba7d832 disk                connected   configured unknown
c1::220203708b8d45f2 disk                connected   configured unknown
c1::220203708b9b20b2 disk                connected   configured unknown
c2                   fc-fabric          connected   unconfigured unknown
c2::50020f2300005f24 disk                connected   unconfigured unknown
c2::50020f2300006107 disk                connected   unconfigured unknown
```

IPFC Management

Note – See also “IPFC Guidelines” on page 9.

To configure IP over Fibre Channel (IPFC), you should have already ensured that the hosts recognize the switch and all attached devices. This chapter describes host recognition of IPFC devices and implementation of IP over Fibre Channel in a SAN. The IPFC driver is based on RFC 2625 and it allows IP traffic to run over Fibre Channel. Configuration of IPFC depends on the instance of the fabric port (fp) driver for the Fibre Channel adapter ports. This chapter contains the following topics:

- “Loading IPFC” on page 86
- “IPFC Invocation, Configuration and Usage” on page 91

Loading IPFC

Configuration of IPFC depends on the instance of fp, or host bus adapter ports. If multiple host bus adapters are present, plumb manually after identifying the fp instance on which IP should be plumbed. The following two procedures determine fp instances.

- “Determining Fibre Channel Adapter Port Instances” on page 86
- “To Plumb an IPFC Instance” on page 89

Determining Fibre Channel Adapter Port Instances

There are two basic ways to determine Fibre Channel adapter port instances to which IP can be plumbed. The first way, which is a short procedure, requires that you know the WWN of the card. In the second way, which is a longer procedure, you do not know the WWN but you do know the physical location of the card. The procedures include:

- “To Determine All Fibre Channel Adapter Port Instances To Which IP Can Be Plumbed With a WWN” on page 87
- “To Determine an FP Instance for a Known HBA Port Physical Device Path” on page 88

▼ To Determine All Fibre Channel Adapter Port Instances To Which IP Can Be Plumbed With a WWN

1. **Become superuser.**
2. **Determine the fp driver instances in your system.**

In the example below, there are four instances (0-3) of fp present in the system.

```
# prtconf -v | grep fp
fp (driver not attached)
fp, instance #0
fp (driver not attached)
fp, instance #1
fp (driver not attached)
fp, instance #2
fp (driver not attached)
fp, instance #3
```

3. **Manually load IPFC to the desired fp instances.**

Use the `ifconfig fcip $interface$ plumb` command, where *interface* is a variable for the desired fp instance number. For example:

```
# ifconfig fcip0 plumb
```

If the command succeeds, the following message appears on the console and messages file:

```
Sep 13 15:52:30 bytownite ip: ip: joining multicasts failed (7) on
fcip1 - will use link layer broadcasts for multicast
```

If no other error message is displayed, manual plumbing has succeeded. Repeat this step for the other fp instances identified in Step 2.

Note – Additional options can be provided to the `ifconfig` command too. See `ifconfig(1M)` for more information. Currently, IPFC can be plumbed on only one port on Fibre Channel adapters with multiple ports. See “IPFC Guidelines” on page 9 for details. An error occurs if a plumb operation is attempted on other ports.

▼ To Determine an FP Instance for a Known HBA Port Physical Device Path

1. Determine the HBA PCI adapter slot and the I/O board PCI slot.

You need this information to perform the calculation in Step 2.

For example, assume you have an array with an HBA card located in PCI adapter slot 5, and the PCI adapter is in slot one of the I/O board.

2. Determine the `fp` instance number.

a. Use an editor to search for the `fp` driver binding name in the

`/etc/path_to_inst` file.

Entries have `fp` on the line.

Note – Determine the correct entry by finding the hardware path according to the server hardware manual or Sun System Handbook. The Sun System Handbook is available at http://sunsolve.sun.com/handbook_pub/.

b. Narrow the search by using the I/O board and slot information from Step 1.

i. Multiply the PCI adapter slot number by the number of adapter ports.

For example, if the HBA has two ports, multiply by two. Using the array with an HBA in the PCI adapter slot five, multiply five by two to get 10.

ii. Add the PCI adapter I/O board slot number to the number derived in Step i.

Using the array with an HBA in PCI adapter slot five and PCI slot one of the I/O board, add one to 10 for a sum of 11.

iii. Convert the number derived in Step ii to hexadecimal.

The number 11 converts to `b` in hexadecimal.

iv. Search for the `fp` entry with `pci@hex` where *hex* is the number you derived in Step iii.

CODE EXAMPLE 6-1 shows a single Fibre Channel network adapter device path. TABLE 6-1 defines the significance of the device path. The instance number for this device path is 7.

CODE EXAMPLE 6-1 PCI Single Fibre Channel Network Adapter Device Path

```
"/pci@b,2000/SUNW,q1lc@2/fp@0,0" 7 "fp"
```


TABLE 6-1 PCI Single Fibre Channel Network Adapter /etc/path_to_inst Device Path Entry

Entry Item	Entry Value
Physical Name	pci@b,2000/SUNW,qlc@2/fp@0.0
Instance Number	7
Driver Binding Name	fp

3. Manually plumb the fp instance.

Use the `ifconfig <interface> plumb` command. In this example, *interface* is `fcip7`.

```
# ifconfig fcip7 plumb
```

If no error message is displayed, manual plumbing has succeeded. Repeat this step for the other `fp` instances identified in Step 2.

▼ To Plumb an IPFC Instance

Use this procedure to load and plumb IPFC.

1. Each FP instance on the system has an entry in `/dev/fc`. If HBAs have been removed, some stale links might exist.

For each entry in `/dev/fc`, issue a `luxadm -e dump_map` command to view all the devices that are visible through that HBA:

```
# luxadm -e dump_map /dev/fc/fp0
Pos  Port_ID Hard_Addr Port WWN          Node WWN          Type
0    610100 0          210000e08b049f53 200000e08b049f53 0x1f
(Unknown Type)
1    620d02 0          210000e08b02c32a 200000e08b02c32a 0x1f
(Unknown Type)
2    620f00 0          210000e08b03eb4b 200000e08b03eb4b 0x1f
(Unknown Type)
3    620e00 0          210100e08b220713 200100e08b220713 0x1f
(Unknown Type,Host Bus Adapter)
# luxadm -e dump_map /dev/fc/fp1
No FC devices found. - /dev/fc/fp1
```

2. Based on the list of devices, determine which destination HBAs are visible to the remote host with which you want to establish IPFC communications.

In the example for this procedure, the destination HBAs have port IDs 610100 and 620d02. The originating HBA's port ID is 620e00.

3. List the physical path of the originating HBA port from which you can see the destination HBA port, where *originating-hba-link* is a variable for the link determined in Step 1:

```
# ls -l /dev/fc/originating-hba-link
```

For example, 0 is the number for the *originating-hba-link*:

```
# ls -l /dev/fc/fp0
lrwxrwxrwx 1 root root          51 Sep  4 08:23 /dev/fc/fp0 ->
../../devices/pci@8,600000/SUNW,qlc@1/fp@0,0:devctl
```

4. Search the physical path identified in Step 3.

You must remove the leading ../../devices from the pathname output. For example

```
# grep pci@8,600000/SUNW,qlc@1/fp@0,0 /etc/path_to_inst
"/pci@8,600000/SUNW,qlc@1/fp@0,0" 0 "fp"
```

5. Determine the *fp* instance for the originating HBA port from the output of the command in Step 4.

The instance number precedes “*fp*” in the output. In the following example output, the instance number is 0.

```
"/pci@8,600000/SUNW,qlc@1/fp@0,0" 0 "fp"
```

6. Use the instance number from Step 5 to load IPFC and plumb the IPFC interface.

In this example, the instance is 0.

```
# ifconfig fcip0 plumb
```

IPFC Invocation, Configuration and Usage

Immediately after installation, start IPFC manually with the `ifconfig` command. On subsequent reboots, the IPFC network interface starts automatically. This section covers:

- “To Start a Network Interface Manually” on page 91
- “To Configure the Host for Automatic Plumbing Upon Reboot” on page 92
- “IPFC Usage” on page 92

▼ To Start a Network Interface Manually

Use this procedure when you want to plumb IPFC with specific netmask values and get the IPFC interface up and running.

1. Use the `ifconfig` command with the appropriate interface.

For example, to enable an IPFC interface associated with `fp` instance 0 and an IP address of `192.9.201.10`, enter:

```
# touch /etc/notrouter
# ifconfig fcip0 inet 192.9.201.10 netmask 255.255.255.0 up
```

The `ifconfig` command is described in more detail in the `ifconfig(1M)` manpage. Ask your network administrator for an appropriate IP address and netmask information.

2. Use the command `ifconfig -a` to verify the network is functioning.

The output of `ifconfig -a` should look like this:

```
lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 index 1
inet 127.0.0.1 netmask ff000000
fcip0: flags=1001843<UP,BROADCAST,RUNNING,MULTICAST,MULTI_BCAST,IPv4> mtu 1500
index 2
inet 192.9.201.10 netmask ffffffff broadcast 192.9.201.255 ether 0:e0:8b:1:3c:f7
hme0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 3
inet 192.9.200.70 netmask ffffffff broadcast 192.9.200.255
ether 8:0:20:fc:e9:49
```

▼ To Configure the Host for Automatic Plumbing Upon Reboot

Each network interface must have an `/etc/hostname.interface` file defining the name of the IP address associated with it. For example, IPFC network *interface* `fcip0` has a file named `/etc/hostname.fcip0`.

1. **Manually create a `/etc/hostname.interface` file with a text editor so it contains a single line that identifies the host name or interface IP address.**
2. **Use a text editor to make any additional entries to the `/etc/inet/hosts` file.**

The Solaris installation program creates the `/etc/inet/hosts` file with minimum entries. You must manually make additional entries with a text editor. (See the `hosts(4)` man page for additional information.)

The `/etc/inet/hosts` file contains the hosts database. This file contains the host names and the primary network interface IP addresses, other network interfaces attached to the system, and any other network interface that the machine must know about.

CODE EXAMPLE 6-2 sun1 machine `/etc/inet/hosts`

```
127.0.0.1      localhost      loghost
192.9.200.70   sun1          #This is the local host name
192.9.201.10   fcip0         #Interface to network 192.9.201.10
```

3. **Edit the `/etc/nsswitch.conf` file so that all uncommented entries have the word `files` before any other nameservice.**

The `/etc/nsswitch.conf` specifies which name service to use for a particular machine.

CODE EXAMPLE 6-3 sun1 machine `/etc/nsswitch.conf` File

```
hosts:      files nis
```

IPFC Usage

Any standard network commands can be used after IPFC is attached. There are not any usage differences when these commands (`telnet`, `ping`, or `ftp`) are used in an Ethernet setup.

Dynamic Reconfiguration On the SAN

When you want to change your system configuration with Dynamic Reconfiguration (DR), device configuration for SAN-based devices must be changed prior to DR. DR is straight-forward with FC-AL devices and only slightly more complicated for fabric devices.

With previously configured FC-AL devices, DR happens automatically upon addition or removal of devices to a host I/O port. With the STMS enabled, the Solaris operating environment host configures the devices as Sun StorEdge Traffic Manager devices. However, DR is more complex with fabric devices. This section covers:

- “DR and Fabric Devices” on page 93
- “To Remove a Fabric Device Before DR” on page 94
- “To Maintain a Fabric Device Configuration With DR” on page 94
- “To Reconfigure Fabric Devices With DR” on page 94
- “To Reconfigure the Sun Enterprise 10000 Server With a Fabric Connection” on page 94

DR and Fabric Devices

Previously configured fabric devices are not automatically reconfigured. Fabric device configurations are not persistent when you remove a system component on which switch-connected host ports reside, and then add that system component back to a host through DR operations. The discussion of on-demand node creation in these sections applies to fabric devices, such as a host port connected to an F port on a switch and an array connected to an F port or TL port on a switch.

▼ To Remove a Fabric Device Before DR

- **Unconfigure the fabric devices that were configured through host ports on the system component with on-demand node creation.**
 - If the STMS is not enabled, see “To Unconfigure a Fabric Device” on page 59.
 - If the STMS is enabled, see “To Unconfigure a Fabric Device Associated With Multipathing Enabled Devices” on page 72.

▼ To Maintain a Fabric Device Configuration With DR

1. **Reconfigure the device through on-demand node creation.**
2. **Perform DR operations according to the instructions in the documentation for the host.**

▼ To Reconfigure Fabric Devices With DR

1. **Add the system component and make it available to the host.**
2. **Reconfigure the device(s) through on-demand node creation.**
 - If the STMS is not enabled, see “To Configure a Fabric Device Without Multipathing Enabled On the Host” on page 55.
 - If the STMS is enabled, see “To Configure Fabric-Connected Device Nodes With Multipathing Enabled Devices” on page 68.

▼ To Reconfigure the Sun Enterprise 10000 Server With a Fabric Connection

The following procedure gives the sequence of operations for a Sun Enterprise 10000 server board with a fabric connection.

1. **Unconfigure the fabric devices on fabric-connected host ports on the board to be detached.**
2. **Start the DR detach operations for the board.**
See the Sun Enterprise 10000 Dynamic Reconfiguration Configuration Guide.
3. **Start the DR attach operations when the board is ready.**
See the Sun Enterprise 10000 Dynamic Reconfiguration Configuration Guide.

4. Configure any fabric devices on the attached boards.

See the sections in Chapter 5 that explain how to recognize the storage devices on the host. On the newly attached board, the devices could be the same or completely new devices.

Glossary

This glossary defines SAN terminology used in this book.

A

- AL_PA** Arbitrated Loop Physical Address. An AL_PA is an eight-bit value used to identify a device attached to a Fibre Channel arbitrated loop.
- Ap_Id** Attachment Point Identifier. The port identifier for either a host or a storage device connected to a switch. For example, `c0` identifies the host port and `c0::50020f23000063a9` identifies an array.

C

- Cascade** Connection of two or more switches together to increase the number of available ports or to allow for increased distances between components of the SAN.

E

- Expansion ports (E port)** Port used to connect FC-SW2 compliant switches to one another.

F

Fabric ports (F port) Fabric port on a Fibre Channel switch. Switch ports used to connect to the fabric capable storage devices (such as the Sun StorEdge T3+ array or Sun StorEdge 39x0, 69x0, 6x20, and 99x0 series), tape drives and host bus adapters.

Fabric One or more switches in a Fibre Channel network. It is also common to refer to something as a "fabric device" or being in "fabric mode."

Fabric Loop ports (FL port) A port that is able to transmit under fabric protocol and also has arbitrated loop capabilities.

FC-SW-2 The second generation of the Fibre Channel Switch Fabric Standard defined by ANSI.

G

Generic port (G port) A port that can operate as either an E port or F port. A port is defined as a G port when it is not yet connected or has not yet assumed a specific function in the fabric.

I

Initiator Each host bus adapter port connection that provides a path to a storage device.

Interswitch Link (ISL) A segment or cable connecting two cascaded switches. Does not include cables from host to the switch or from storage devices to the switch.

N

N port A Fibre Channel port that supports point-to-point or fabric connections.

Name Server Zones (NS zones)

A set of name server (NS) ports. The ports receive name server information (port number, type, address, WWN, and so on). NS zones can be port-based or WWN-based.

P

Private Loop

Fibre Channel Arbitrated Loop (FC-AL) with 8-bit addressing that supports up to 126 device connections with no fabric attachment (F and FL ports). Private loops are closed systems incapable of seeing outside the loop.

Public Loop

Arbitrated loop that supports fabric login and services. Provides 24-bit Fibre Channel addressing and up to 16 million node connections fabric wide. Uses name server (NS) ports.

S

Segmented Loop ports (SL ports)

A port connected to a private loop device. SL ports make a switch behave like a hub, but with the advantage of better performance and the ability to segment the private loop into SL Zones for ease of administration and isolation of resources. Ports in SL zones do not communicate with ports in NS zones. Supports the Sun StorEdge A5200, A3500FC arrays and FC tape devices. SL zones contain SL ports only. SL ports are not supported in the SAN 4.x release but were in earlier releases.

Segmented Loop Zones (SL zones)

A set of SL ports on the switch that behave as a single private loop. This grouping behaves like a hub, but with the advantage of better performance and the ability to segment the private loop into SL zones for ease of administration and isolation of resources. SL zones are not supported in the SAN 4.x release but were in earlier releases.

T

Transfer or Trunk ports (T ports)

Switch ports from the SAN 3.0 release used to connect to other switches in a cascade. T ports are used to attach a port on one switch to a port on another switch. Replaced by E ports in SAN 4.0 release.

Translative Loop ports (TL ports)

Ports on the switch that allow private-to-public and public-to-private loop connectivity. TL ports present private devices to a fabric as if they were public so they can communicate with the off-loop devices, such as fabric devices connected to F ports and FL ports and private devices on other TL port loops. Supports the Sun StorEdge T3 arrays. Hosts recognize arrays that use TL ports as fabric devices.

Z

Zone A set of ports or WWNs and their connected devices that have been grouped together to control information exchange.

Zoning Function of the switch that allows segregation of devices. Zoning is done for a variety of reasons, such as security, simplicity, performance, or dedication of resources.

Index

SYMBOLS

/etc/ethers file, 12
/etc/hosts file, 13
/etc/nsswitch.conf file, 13

NUMERICS

24-bit addressing devices, 44
24-bit Fibre Channel addressing devices, 51

A

Ap_Id, 74
array port, 44

B

boot block, 35
boot disk, 24
boot disk change, 29
boot disk file systems, 34, 35, 38
boot disk labeling, 33
boot disk partitioning, 25
boot disks as boot devices, 39
boot image, 19
boot image modification, 21
boot server setup, 20
bootability considerations, 16

bootability process overview, 19
booting, 15
broadcasting, 9

C

cascading
 guidelines, 8
CD 0, 16
cfgadm
 -al, 53
 -al -o show_FCP_dev, 54
 -c configure, 46
 -c unconfigure, 49
 -fp, 65
 -l, 53
cfgadm -al, 61, 75
cfgadm -al -o show_FCP_dev, 62
cfgadm -c configure, 61
cfgadm(1M), 52
client setup for bootability, 22
configuration, 11
 fabric devices, 52

D

device
 activity, 72
 configuration, 52
 node creation, 46

- node discovery, 61
- node identification, 54
- node removal, 52
- removal, 73
- storage, 51
- device node
 - creation, 52
 - removal, 61
- device node creation, 61, 65
- device node removal, 65
- device path name, 29
- drivers, 54
- Dump and Restore boot method, 17
- Dump and Restore bootability method, 24

E

- E_port, 100
- eight-bit addressing devices, 44
- Ethernet port, 11

F

- F_port, 44
- fabric boot, 16, 24
- fabric connectivity, 44
- fabric-connected host ports, 53
- FC-AL, 99
 - devices, 51
- fc-fabric, 68
- FCode level, 16
- Fibre Channel
 - Fibre Channel Arbitrated Loop, 99
 - Protocol, 54
- file system creation on boot disks, 34
- file systems on boot disks, 38
- format command, 26

G

- guidelines
 - cascading, 8

H

- HBA, 15
- high availability, 10
- host bus adapter, 15, 51
- host port, 44
- hub-to-switch connectivity, 8

I

- initiator, 98
 - active, 98
 - inactive, 98
- install image, 19
- install server setup, 20
- installboot command, 35
- inter-switch link, 98
- IP address, 11
- IPFC
 - guidelines, 9
 - SCSI LUN level information, 54
- ISL, 98

J

- Jump Start server, 17

L

- labeling the boot disk, 33
- load balancing, 10
- loading drivers, 54
- long-wave
 - GBIC, 8
- LUN, 10
 - level information, 54
 - recognition, 46
- luxadm, 56, 75
- luxadm qlgc_s_download, 16

M

- manual boot, 22

- mini-root, 19
- modinfo, 54
- multicasting, 9
- multipathing, 10, 65
 - devices and configuration, 68
 - utilities, 45
- multiple device node, 70, 81
 - creation, 77
 - removal, 77

N

- name server
 - zones, 99
- NAS, 9
- Net Install Patching boot method, 17
- NFS, 9
- non-fabric boot, 17, 24
- non-fabric device, 17
- non-fabric devices, 17
- NS
 - zones, 99

P

- packages, 23
- partition command, 27
- partition ID tag, 31
- partition layout, 26, 29
- partition permission flags, 32
- partition size, 32
- partition table, 27
- partitioning, 25
- patchadd command, 20, 25
- patches, 23
- physical device representation, 69
- pkgadd command, 20, 25
- port, 44
 - N_Port, 98
 - segmented loop, 99
 - transfer, 100
 - trunk, 100
 - WWN listing, 74
- port list, 67

- port-based zones, 10, 99
- print command, 27
- private loop, 99
 - devices, 51
- promiscuous mode, 9
- public loop, 99

Q

- qlc number, 29

R

- RARP, 11
- Reverse Address Resolution Protocol, 11
- root file system, 35

S

- Segmented Loop
 - zones, 99
- setup_install_server command, 20
- short-wave
 - GBIC, 8
 - tranceiver, 8
- show_FCP_dev, 56, 62
- shutdown command, 22
- SL
 - zones, 99
- SL_Ports, 99
- slices, 26, 28, 31
- SNDR, 9
- snoop, 9
- Solaris installation, 19
- Solaris OE documentation, 35
- ssd driver, 54
- st driver, 54
- STMS, 24
- storage devices, 49
- StorEdge T3+ array, 43
- Sun StorEdge 2G FC PCI Dual Channel Network Adapter, 16

- Sun StorEdge 2G FC PCI Single Channel Network Adapter, 16
- Sun StorEdge 39x0 series, 43, 44
- Sun StorEdge 69x0 series, 43, 44
- Sun StorEdge 99x0 series, 43, 44
- Sun StorEdge T3 device nodes, 67
- Sun StorEdge Traffic Manager tool, 66
- Sun switch, 11
- suninstall program, 22
- support contact, xiv
- switch
 - switch-to-switch connection, 8
- System Administration Guide
 - Basic Administration for Solaris 9 OE, 35
- System Administration Guide, Volume 1 for Solaris 8 OE, 35

T

- T_Ports, 100
- temporary boot disk, 24, 37
- TL
 - port, 44, 100
- Tools directory, 20
- transfer ports, 100
- translation loop
 - port, 44
- translative
 - loop port, 100
- trunk ports, 100

U

- ufsdump command, 35
- ufsrestore command, 35
- unconfiguring
 - a single path, 75
 - devices, 59
 - multipathed devices, 74
 - multiple devices, 83
- unconfiguring devices, 73

V

- vfstab file, 37
- volumes, 10

W

- WWN-based zones, 99
- WWNs, 100

Z

- zone, 100
 - name server, 99
 - segmented loop, 99
- zoning, 6, 100