

FCP Solaris™ Host Utilities 4.2 for Veritas®

Installation and Setup Guide

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number 215-03235_A0
July 2007

Copyright and trademark information

Copyright information

Copyright © 1994–2007 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted Network Appliance material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETWORK APPLIANCE “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL NETWORK APPLIANCE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, Data ONTAP, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, SyncMirror, Topio, VFM, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. Cryptainer, Cryptoshred, Datafort, and Decru are registered trademarks, and Lifetime Key Management and OpenKey are trademarks, of Decru, a Network Appliance, Inc. company, in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexShare, FPolicy, HyperSAN, InfoFabric, LockVault, NOW, NOW NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simplicore, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, SohoFiler, SpinMirror, SpinRestore, SpinShot, SpinStor,

StoreVault, vFiler, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the U.S.A.

IBM, the IBM logo, AIX, and System Storage are trademarks and/or registered trademarks of International Business Machines Corporation.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the U.S.A. and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	Host Utilities Installation and Setup Overview	1
	Host Utilities Overview.	2
	FCP Solaris Host Utilities installation overview	5
	Where to go for more information	9
Chapter 2	Installing Emulex Software	11
	Before you begin	12
	Installing HBA software	13
	Upgrading firmware	18
Chapter 3	Configuring the Solaris Host	21
	Installing the FCP Solaris Host Utilities	22
	Completing the installation and configuring system parameters.	25
	Setting parameters manually	26
	Using the basic_config command	28
	Removing an FCP Solaris Host Utilities or Attach Kit installation	31
Chapter 4	Completing the Configuration	33
	Configuration overview.	34
	Mapping igroups and LUNs	36
	Creating persistent bindings	37
	Configuring the host to discover LUNs	40
	Labeling the LUN and configuring Volume Management software	42
Appendix A	Using the create_binding.pl Script	45
	The create_binding.pl script	46

	Before you begin	48
	Using the create_binding.pl script	50
Appendix B	Configuring SAN Boot	55
	SAN boot overview.	56
	Configuration overview.	57
	Setting up the Emulex HBA	59
	Creating the bootable LUN.	66
	Copy bootblks and data.	72
	Modifying OpenBoot.	75
Appendix C	Changing cfmode to Single Image	81
	Understanding cfmode	82
	Completing prerequisites	84
	Upgrading to single image mode	85
Appendix D	Veritas and Array Support Library	91
	Overview of the Array Support Library	92
	Installing the ASL	93
	Removing the ASL	96
	Obtaining multipathing information	97
	Optimizing the Veritas DMP Restore Daemon.	101
	Understanding ASL error message types.	102
Appendix E	Using sanlun	103
	Displaying host LUN information	104
	Displaying information about the storage system paths	106
	Displaying host HBA information	108
Appendix F	Tips and Troubleshooting.	111

Confirming that your system is correctly set up 112

Using Diagnostic Utilities 118

Collecting storage system information 119

Collecting Solaris information 121

Collecting switch information 122

Index 125

Preface

About this guide

This guide describes how to configure a Solaris™ host in a SAN environment. The host accesses data on a NetApp® storage system that runs Data ONTAP® software with the Fibre Channel Protocol (FCP).

See the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products*, available on the NetApp on the Web™ (NOW) site at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/ for information on the specific Solaris versions supported.

Audience

This guide is for system installers and administrators who:

- ◆ Are familiar with the Sun Solaris operating system
- ◆ Understand how to configure and administer a storage system
- ◆ Understand how to use FCP to store and retrieve data on a storage system.

Terminology

This guide uses the following terms:

- ◆ NetApp storage products (filers, FAS appliances, and NearStore systems) are all storage systems and are also sometimes called filers or storage appliances.
- ◆ *igroup* refers to a collection of unique identifiers, either FCP WWPNs in a SCSI network or iSCSI node names of initiators (hosts) in an IP network, that are given access to LUNs when they are mapped to those LUNs.
- ◆ *LUN* (Logical Unit Number) refers to a logical unit of storage identified by a number.
- ◆ *LUN ID* refers to the numerical identifier for a LUN.
- ◆ *Initiator* refers to the system component that originates an I/O command over an I/O bus or network.
- ◆ *Target* refers to the storage system that receives an FCP I/O request and serves the data requested by the initiator.

Command conventions

You can enter storage system commands on the system console or from any client that can obtain access to the storage system using a Telnet session.

Formatting conventions

The following table lists different character formats used in this guide to set off special information.

Formatting convention	Type of information
<i>Italic type</i>	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fc test adaptername</code> command, you enter the characters “fc test” followed by the actual name of the adapter. ◆ Book titles in cross-references.
Monospaced font	◆ Command and daemon names. ◆ Information displayed on the system console or other computer monitors. ◆ The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

Keyboard conventions

This guide uses capitalization and some abbreviations to refer to the keys on the keyboard. The keys on your keyboard might not be labeled exactly as they are in this guide.

What is in this guide...	What it means...
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
<i>Enter, enter</i>	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. Also used to mean pressing one or more keys and then pressing the Enter key.
<i>type</i>	Used to mean pressing one or more keys on the keyboard.

Special messages

This guide contains special messages that are described as follows:

Note

A note contains important information that helps you install or operate the system efficiently.

Caution

An attention contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

How to send your comments

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this publication, you can send us comments electronically to doccomments@netapp.com. To help us direct the comments to the correct division, please put FCP Solaris in the subject line.

About this chapter

This chapter provides a brief overview of the FCP Solaris Host Utilities for Veritas. The Veritas for this release refers to Veritas Storage Foundation 5.0.

In addition, this chapter provides a high-level overview of the tasks you need to perform to install the Host Utilities and set up your system. It also includes pointers to other documentation that might be helpful in working with your Host Utilities setup.

Topics covered in this chapter

This chapter includes the following topics:

- ◆ [“Host Utilities Overview”](#) on page 2
- ◆ [“FCP Solaris Host Utilities installation overview”](#) on page 5
- ◆ [“Where to go for more information”](#) on page 9

Host Utilities Overview

Important information about Host Utilities and Attach Kits

Previously NetApp provided FCP Solaris Attach Kits to enable users to connect Solaris hosts to NetApp storage systems. Going forward, NetApp is using Host Utilities to provide this support.

Host Utilities both simplify the products and provide more options for the host stacks. A single FCP Solaris Host Utilities can contain software tools and documentation packages to support different host configurations.

Because the Host Utilities are the umbrella for the different kits, component version numbers are not tied to the Host Utilities version number. It is possible for each component to have a different version number. For example, at the time this document was produced, it supported the FCP Host Utilities 4.2 for Veritas using version 3.4 of the SAN Toolkit.

If you are concerned about the version numbers, you can confirm that you have the most current components by checking the Host Utilities version number in the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Overview of the FCP Solaris Host Utilities

The Host Utilities provide the tools and information necessary to enable a Solaris host to connect to your storage systems. The Host Utilities consist of:

- ◆ The SAN Toolkit, which contains:
 - ❖ Configuration tools that help you configure host, HBA, and system files and create persistent bindings.
You can configure the Solaris host using the tools provided in the FCP Solaris Host Utilities, HBA provided tools, or a combination FCP Solaris Host Utilities and HBA provided tools.
 - ❖ The `san lun` utility that helps you manage LUNs and the host HBA.
 - ❖ Diagnostic scripts that provide diagnostic information about components in your configuration. Customer Support may ask for output from these scripts.
- ◆ Documentation that describes how to configure the Solaris host. It includes this setup guide, *Release Notes*, and a quick command reference guide.

The most current information on the supported stack is in the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>. NetApp adds components on an ongoing basis. This online site is updated frequently to provide you with information about the current stack.

Software tools provided by the Host Utilities

The FCP Solaris Host Utilities bundle the tools into a SAN Toolkit. The toolkit contains the following components:

- ◆ `basic_config` command line tool. This tool modifies the Fibre Channel tunable variables to set them to values required by the FCP driver.
- ◆ `filer_info` script. This script collects information about the storage system.
- ◆ `mpxio_set` script. This script lets you clean up the I/O connection if you were running the iSCSI Solaris Support Kit 1.0. (If you were not using the support kit, you should not need to use this script.)
- ◆ `san_version` command. This command displays the version of the SAN Toolkit that you are running.
- ◆ `sanlun` utility. You can use this utility to display information about LUNs on the storage system that are available to this host.
- ◆ `solaris_info` script. This script collects configuration information about your host.
- ◆ Additional diagnostic scripts. The following scripts can also be used to collect troubleshooting information:
 - ❖ `brocade_info` collects information about Brocade switches installed in the network.
 - ❖ `cisco_info` collects information about Cisco switches installed in the network.
 - ❖ `mcdata_info` collects information about McData switches installed in the network.
 - ❖ `qlogic_info` collects information about QLogic switches installed in the network.
 - ❖ `SHsupport.pm` is the support library for the *info scripts.
 - ❖ `Telnet.pm` support library for *info scripts.
 - ❖ `version` displays the version of the `sanlun` command.
- ◆ The man pages for `sanlun` and the diagnostic scripts.

Remove previous versions before installing Host Utilities

The FCP Solaris Host Utilities 4.2 for Veritas uses a different stack from previous versions of the Host Utilities. You should not have both products installed on the same host.

If you have an earlier version installed, you must remove it before you install this version.

- ◆ See “[Removing an FCP Solaris Host Utilities or Attach Kit installation](#)” on page 31 for additional information.

FCP Solaris Host Utilities installation overview

Overview of Installing the Host Utilities and setting up your system

Installing the FCP Solaris Host Utilities and setting up your system involves a number of steps that are performed on both the storage system and the host. The following sections provide information about the tasks you need to perform.

Pre-requisites

Before you install the Host Utilities, you must perform the following tasks:

1. Set up your system:
 - ❖ Host OS and appropriate updates
 - ❖ HBAs
 - ❖ Drivers
 - ❖ Storage system and Data ONTAP
 - ❖ Switches (if used)

See the *Compatibility and Configuration Guide for NetApp FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml> for current information about the system requirements.

2. Verify that your storage system is:
 - ❖ Licensed correctly
 - ❖ Using the correct cfmode
 - ❖ Running the FCP service
 - ❖ Set up to use Veritas Volume Manager
 - ❖ Configured to work with the target HBAs
 - ❖ Set up to work with the Solaris host and the initiator HBAs
 - ❖ Set up with working volumes and qtrees (if desired)

Note

It is recommended that you create volumes and LUNs with space reservations enabled.

3. If you're using a fabric connection, verify that the switch is:
 - ❖ Set up correctly
 - ❖ Zoned (if appropriate for your system)

- ❖ Cabled according to the instructions that came with the switch
4. Verify that:
- ❖ The power to the switches, disk shelves, storage systems, and hosts is turned on in the correct order
 - ❖ The host and the storage system can communicate

Installation overview

Installing the Host Utilities involves the following tasks:

1. Get a copy of the Host Utilities' SAN Toolkit software package.
 - ❖ Get the compressed file containing the software package by either downloading it or copying it from the CD.
 - ❖ Extract the software package from the compressed file.
2. Install the toolkit package:
 - ❖ Go to the directory containing the extracted software package.
 - ❖ Use the command `pkgadd -d NTAPSANTool.pkg` to install the toolkit.
 - ❖ Configure storage specific variables in the `/kernel/drv/lpfc.conf` and `/etc/system` files. You can use the `basic_config` command to set the driver parameters.
3. Configure the LUNs by performing the following tasks:
 - ❖ Create at least one igroup and at least one LUN.
One way to create igroups and LUNs is to use the `lun setup` command to create igroups and LUNs.
 - ❖ Map the LUN to the igroup.
 - ❖ Creating persistent bindings using `lputil`. (You can also create persistent bindings using the NetApp provided `create_binding.pl` script.)
 - ❖ Configure the host to discover the LUNs.
In the `sd.conf` file, create an entry for every NetAppLUN.
 - ❖ Label the LUNs using the Solaris format utility
Use the `/usr/sbin/format` utility.
 - ❖ Configure the Volume Management software
 - ❖ Display information about the LUNs and HBA.
Use the `sanlun` command.

For detailed information on these tasks

The following sections in this document provide more details about setting up your system, installing and configuring the Host Utilities, and creating igroups and LUNs:

- ◆ [“Installing Emulex Software”](#) on page 11
- ◆ [“Confirming that your system is correctly set up”](#) on page 112
- ◆ [“Installing the FCP Solaris Host Utilities”](#) on page 22
- ◆ [“Completing the installation and configuring system parameters”](#) on page 25
- ◆ [“Mapping igroups and LUNs”](#) on page 36
- ◆ [“Configuring the host to discover LUNs”](#) on page 40
- ◆ [“Labeling the new LUN”](#) on page 42
- ◆ [“Configuring Volume Management”](#) on page 43
- ◆ [“Displaying host LUN information”](#) on page 104
- ◆ [“Displaying information about the storage system paths”](#) on page 106
- ◆ [“Displaying host HBA information”](#) on page 108

In addition, the section [“Where to go for more information”](#) on page 9 provides pointers to additional documents and web sites that can help you with your system setup.

Optional procedures

You may wish to complete the following tasks after installation is complete.

Task	Additional information
Use the <code>create_binding.pl</code> script to create persistent bindings. This script is bundled with the FCP Solaris Host Utilities software.	See Appendix A, “Using the create_binding.pl Script,” on page 45
Configure SAN boot. SAN boot uses a SAN attached disk, like a LUN, as a boot device for a Solaris host	See Appendix B, “Configuring SAN Boot,” on page 55
Update the storage system cluster setting to single image mode.	See Appendix C, “Changing cfmode to Single Image,” on page 81.
Install the Array Support Library (ASL)	See Appendix D, “Veritas and Array Support Library,” on page 91

Task	Additional information
Use sanlun.	See Appendix E, “ Using sanlun ,” on page 103
Run diagnostic scripts. Normally, you would only run these scripts if you encountered a problem. These scripts collect information about your system and place it in compressed file that you can send to customer support.	See Appendix F, “ Using Diagnostic Utilities ,” on page 118

Where to go for more information

Where to go for host and storage system information

The following table describes the documents on the NetApp on the Web (NOW) site at the time this *Installation and Setup Guide* was written. Use the links below to access information about host and storage system requirements, supported configurations, and cabling instructions.

If you need information about...	See the following NetApp product documentation...
Known issues, system requirements, and last minute updates	The <i>Release Notes</i> , which are available at http://now.netapp.com/NOW/knowledge/docs/san/
A summary of some of the commands you might use with the Host Utilities	The <i>Quick Reference</i> for the FCP Solaris Host Utilities, which is available from the NOW site
The latest supported configurations	◆ <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml ◆ <i>System Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Supported Fibre Channel SAN topologies	<i>FCP Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/
Configuring the storage system	◆ <i>Best Practices for Reliability: New System Installation</i> at http://now.netapp.com/NOW/knowledge/docs/other/best_practices/index.shtml ◆ <i>Data ONTAP® Software Setup Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>Data ONTAP® Block Access Management Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml

If you need information about...	See the following NetApp product documentation...
Managing SAN storage on the storage system	◆ <i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>Data ONTAP® Release Notes</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>Command Reference</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>FilerView® Online Help</i> - available from within the FilerView GUI ◆ <i>Best Practices for Reliability: Ongoing Maintenance</i> at http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm ◆
Verifying compatibility of storage system with environmental requirements	◆ <i>Site Requirements</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Troubleshooting, operational considerations, and post-release information	◆ FCP Solaris Host Utilities Release Notes at http://now.netapp.com/NOW/knowledge/docs/san/
Upgrading Data ONTAP	◆ <i>Data ONTAP® Upgrade Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml
Installing and configuring the HBA in your host	◆ Refer to your HBA vendor documentation.
Your host operating system	◆ Refer to your operating system documentation. You can download Sun manuals in PDF format from http://docs.sun.com/.
Veritas Volume Manager	◆ Refer to your Veritas documentation.
Working with Emulex	◆ Refer to the Emulex documentation. The Emulex partner site for NetApp contains links to the documentation; see http://www.emulex.com/netapp/support/index.jsp

Configuration overview

This chapter explains how to obtain and install the Emulex software, including the Emulex utility programs, which are provided as part of the emlxu utilities kit. It also describes how to upgrade firmware.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Before you begin”](#) on page 12
- ◆ [“Installing HBA software”](#) on page 13
- ◆ [“Upgrading firmware”](#) on page 18

Before you begin

Complete the pre-requisite tasks

This manual assumes that you have installed and configured your storage system, connected it to the switch or fabric, and installed and connected the HBA on the host.

The following describes the tasks you should have completed before you download and configure the Emulex software.

Configuration Task	Additional information
Configure the storage system ◆ Install Data ONTAP ◆ Install licenses ◆ Connect the storage system to the fabric or switch	◆ Index to Data ONTAP documentation at http://now.netapp.com/NOW/knowledge/docs/ontap ◆ The <i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ Switch documentation at http://now.netapp.com/NOW/knowledge/docs/san/
Configure the Solaris host ◆ Verify OS version and patches for the FCP Solaris Host Utilities ◆ Install HBA hardware ◆ Verify cables to fabric or switch	◆ The <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml ◆ HBA documentation provided with the HBA you use ◆ Index to switch and cable documentation at http://now.netapp.com/NOW/knowledge/docs/san/

Installing HBA software

Obtaining the software

The Emulex software and documentation is available on the NetApp partner page. Go to <http://www.emulex.com/ts/> and select NetApp.

If you are planning to use SAN boot (see “[Configuring SAN Boot](#)” on page 55), you must also download and install the EMLXemlxu utilities kit. The SAN boot procedure requires that you use the EMLXemlxu utility and the EMLXdrvdriver.

Emulex components

The Emulex software is a .tar file that includes the following software components:

- ◆ The Host Bus Adapter driver software
- ◆ The HBA management suit, which includes HBAnyware and lputil.

This document uses lputil to install new firmware and to create persistent bindings. You can also perform these tasks using HBAnyware. If you prefer HBAnyware, refer to Emulex documentation for configuration information.

Downloading Emulex components

The latest HBA firmware is available on the NetApp page of <http://www.emulex.com/netapp/support/index.jsp>. Because your HBA may use a previous firmware version, and you may wish to download new firmware at the same time as you download the rest of the Emulex software.

Follow these steps to download and extract Emulex software and firmware.

Step	Action
1	On the Solaris host, create a download directory for the Emulex software. Enter: mkdir /tmp/emulex

Step	Action
2	To download Emulex driver and firmware, go to http://www.emulex.com/ts/ and click the NetApp link. 1. Follow directions to download the driver software to the /tmp/emulex directory. 2. Follow directions to download Emulex firmware to the /tmp/emulex directory.
3	Change to the /tmp/emulex directory and extract the .tar files that contain the Emulex software. Enter: <pre>tar -xvf <solaris_version>-<driver_version>.tar</pre> Result: The .tar files are extracted to the default directory.
4	Extract the .tar files that contain Emulex applications and driver software. Example: The following commands extract Emulex driver and management software. <pre>tar -xvf <EmLxApps><version_number>-solaris.tar tar -xvf lpfc-<version_number>-sparc.tar</pre>
5	Unzip and extract the HBAnyware software. Example: The following commands unzip and extract HBAnyware software: <pre>gzip -d HBAnyware-<version_number>-sparc.tar.gz tar -xvf HBAnyware-<version_number>-sparc.tar</pre>

Install the EMLXemlxu utilities

Download the EMLXemlxu utilities kit from Emulex Web site. This kit contains the Emulex Fibre Channel Adapter (FCA) utilities. You use tools in the utilities to unbind the Emulex HBA and to set up SAN boot. To install the utilities, follow the instructions in the Emulex *Installing and using the FCA utilities* manual.

Note

At the time this document was produced, the link to the utilities was <http://www-dl.emulex.com/support/solaris/602h/emulex.html>.

Unbinding the HBA Before you continue, unbind the Emulex HBA from the emlxs driver. Complete the following steps.

Step	Action
1	Start the <code>emlxdrv</code> tool, which is in <code>/opt/EMLXemlxu/bin/emlxdrv</code>. When you start this tool, it displays information about the drivers, which are using emlxs names. Note _____ By default, when the SFS is enabled, it gives the emlxs names to the Emulex HBA devices. Example: Executing the <code>emlxdrv</code> tool checks the device names and displays a list of commands available with the tool. <pre> #/opt/EMLXemlxu/bin/emlxdrv EMLXDRV Driver Management Utility, Version 1.00j COPYRIGHT © 2004-2005 Emulex. All rights reserved. Driver Alias Present Boot Sun Models ----- emlxs lpfs no no no LP8000S and LP9002S (SBUS) emlxs f800 no no no LP8000 and LP8000DC emlxs f900 no no no LP9002, LP9002C, LP9002DC, and LP9402DC emlxs f980 no no no LP9802 and LP9802DC emlxs fa00 yes no no LP10000, LP10000DC and LP10000ExDC emlxs fd00 no no no LP11000 and LP11002 emlxs fe00 no no no LPe11000 and LPe11002 emlxs f0a5 no no no 2G Blade Adapter (emlxs only) emlxs fc00 no no yes LP10000-S and LP10000DC-S emlxs fc20 no no yes LPe11000-S and LPe11002-S Available commands: ----- set_emlxs <Alias> - Sets emlxs driver to bind to the specified device(s) set_emlxs_sun - Sets emlxs driver to bind to all Sun devices set_emlxs_all - Sets emlxs driver to bind to all devices set_lpfc <Alias> - Sets lpfc driver to bind to the specified device(s) set_lpfc_nonsun - Sets lpfc driver to bind to all non-Sun devices clear_dev <Alias> - Clears driver binding to the specified device(s) clear_lpfc - Clears all lpfc driver bindings clear_emlxs - Clears all emlxs driver bindings clear_sun - Clears driver bindings to all Sun devices clear_nonsun - Clears driver bindings to all non-Sun devices clear_all - Clears driver bindings to all devices q - Exits this program.</pre>

Step	Action
2	Enter the <code>clear_emlxs</code> command to remove the emlxs driver bindings from the HBA. Example: Executing the <code>clear_emlxs</code> command removes the bindings. As you can see from the output of the command, the Emulex HBA no longer uses emlxs names. <pre># emlxdrv> clear_emlxs Driver Alias Present Boot Sun Models ----- - lpfs no no no LP8000S and LP9002S (SBUS) - f800 no no no LP8000 and LP8000DC - f900 no no no LP9002, LP9002C, LP9002DC, and LP9402DC - f980 no no no LP9802 and LP9802DC - fa00 yes no no LP10000, LP10000DC and LP10000ExDC - fd00 no no no LP11000 and LP11002 - fe00 no no no LPe11000 and LPe11002 - f0a5 no no no 2G Blade Adapter (emlxs only) - fc00 no no yes LP10000-S and LP10000DC-S - fc20 no no yes LPe11000-S and LPe11002-S</pre>
3	Enter <code>q</code> to quit the emlxdrv tool.

Adding the Emulex packages

To add the Emulex packages, follow the steps below.

Step	Action
1	From the directory where you downloaded the Emulex software, enter: <pre>pkgadd -d .</pre>
2	At the Select packages you wish to process prompt, enter the numbers that correspond to the Emulex driver and driver management software. Result: The system displays the pkgadd installation script.

Step	Action
3	Respond as follows to the pkgadd prompts: 1. Rebuild manual pages database for section 7d [y,n,?] y 2. Use IP networking over Fibre Channel [y,n?] n This option is not supported on FCP-attached NetApp storage systems. 3. Do you want to continue with the installation of <lpfc> [y,n,?] y 4. Do you want to continue with the installation of <HBAnyware> [y,n,?] y 5. Select package(s) you wish to process (or "all" to process all packages). (default:all) [?,??,q] q
4	Depending on your configuration, take one of the actions below. ◆ Update HBA firmware, if required —follow the instructions in “Firmware upgrade procedure” on page 19 ◆ Configure the host system — follow the procedures in Chapter 3, “Configuring the Solaris Host,” on page 21 Note_____ You must reboot the host with the reconfigure option (reboot -- -r). _____ For an example of the type of output you see after the lpfc driver has been bound to the HBA, see the example in Step 9 on page 65.

Upgrading firmware

Determining the firmware version

Follow the steps below to determine if you need to upgrade Emulex firmware.

Step	Action
1	Activate the lputil tool. Enter: <code>/usr/sbin/lpfc/lputil</code> Result: The software displays the Main Menu.
2	At the Enter choice prompt that follows the Main Menu, enter 2 to display adapter information. Result: The software displays the Adapter Information Menu.
3	At the Enter choice prompt that follows the Adapter Information Menu, enter 2 to display adapter revision information. Result: The software displays a choice of adapters.
4	At the Enter choice prompt that follows adapter choices, select the adapter for which you want revision information. Result: The software displays the firmware version for the adapter you specify.
5	To exit lputil: 1. Enter 0 to continue. 2. At the Enter choice prompt after the Adapter Information, enter 0 to return to the Main Menu. 3. At the Enter choice prompt after the Main Menu, enter 0 to exit lputil.
6	Access the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products Guide</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config to determine the current firmware requirement.

Firmware upgrade procedure

Follow the steps below to upgrade your firmware version.

Note

You can always download the latest firmware from the Network Appliance page on www.emulex.com/ts/index.htm.

Step	Action
1	Activate the lputil tool. Enter: <code>/usr/sbin/lpfc/lputil</code> Result: The software displays the Main Menu.
2	At the Enter choice prompt that follows the Main Menu, enter 3 to being the firmware upgrade. Result: The software displays the Firmware Maintenance Menu and prompts you for the hardware adapter.
3	At the Enter choice prompt, select the adapter for which you want to upgrade firmware. Result: The software displays the Firmware Maintenance Menu.
4	At the Enter choice prompt, select 1 to load the firmware image. Result: The software prompts you for the location of the Emulex firmware.
5	Enter the specification for the directory where you downloaded the Emulex firmware. Result: The software installs the firmware and displays the Firmware Maintenance Menu.
6	To exit lputil: 1. At the Enter choice prompt after the Firmware Maintenance Menu, enter 0 to return to the Main Menu. 2. At the Enter choice prompt after the Main Menu, enter 0 to exit lputil.

Chapter overview

This chapter explains how install the FCP Solaris Host Utilities software and configure the Solaris host.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Installing the FCP Solaris Host Utilities”](#) on page 22
- ◆ [“Completing the installation and configuring system parameters”](#) on page 25
- ◆ [“Setting parameters manually”](#) on page 26
- ◆ [“Using the basic_config command”](#) on page 28
- ◆ [“Removing an FCP Solaris Host Utilities or Attach Kit installation”](#) on page 31

Installing the FCP Solaris Host Utilities

Key steps involved in setting up the Host Utilities

To set up the Host Utilities, you must install the Host Utilities SAN Toolkit software. Once you've installed it, you can use the tools it provides to configure your host parameters. The Host Utilities provide the following configuration tools:

- ◆ `basic_config` configures the host system parameters
- ◆ `create_binding.pl` creates persistent bindings
- ◆ `sanlun` displays host and storage system information

You can also use tools from the HBA Management Suite to perform these functions. If you wish to use HBA management suite tools, you do not need to install the FCP Solaris Host Utilities. To configure the host using HBA management tools, refer to “[Completing the installation and configuring system parameters](#)” on page 25.

Before you begin

Before you install the Host Utilities SAN Toolkit software package:

- ◆ You must verify that your host system meets requirements before you begin the installation. Refer to the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/ for the latest hardware and software requirements.
- ◆ If you currently have the FCP Solaris Host Utilities 4.0 or FCP Solaris Attach Kit installed, you must remove it. See “[Removing an FCP Solaris Host Utilities or Attach Kit installation](#)” on page 31.

Key steps involved in setting up the Host Utilities

Once your host is set up, you must install the Host Utilities SAN Toolkit software package. Once you've installed the software, you can use the tools it provides to configure your host parameters. The Host Utilities provide the following configuration tools:

- ◆ `basic_config` configures the host system parameters
- ◆ `create_binding.pl` creates persistent bindings
- ◆ `sanlun` displays host and storage system information

You can also use tools from the HBA Management Suite to perform these functions. If you wish to use HBA management suite tools, you do not need to install the FCP Solaris Host Utilities. To configure the host using HBA management tools, refer to “[Completing the installation and configuring system parameters](#)” on page 25.

Obtaining the Host Utilities software

The software is provided on the CD-ROM in your Host Utilities media kit. You can also download the software from the NOW web site at <http://now.netapp.com/NOW/cgi-bin/software>.

Installing the Host Utilities from NOW

To install the FCP Solaris Host Utilities software from NOW, complete the following steps.

Note

NetApp recommends that you do not install the FCP Solaris Host Utilities 4.2 for Veritas on systems using versions of the Veritas Storage Foundation prior to 5.0.

Step	Action
1	Download or copy the <code>santoolkit_solaris_sparc_3.4.tar.Z</code> file to a directory on your host. Enter: <pre>cp santoolkit_solaris_sparc_v3.4.tar.Z /<directory_name></pre> Note The actual file name for the Host Utilities software may be slightly different from the one shown here.
2	Unzip the file. Enter: <pre>gunzip santoolkit_solaris_sparc_v3.4.tar.Z</pre> Result: The software unzips the tar.Z files
3	Untar the file. Change to the directory where you copied the file, and enter: <pre>tar -xvf santoolkit_solaris_sparc_v3.4.tar.Z</pre> Result: The Host Utilities scripts are extracted to the default directory.

Step	Action
4	Add the packages that you extracted from tar file. Enter: <pre>pkgadd -d NTAPSANTool.pkg</pre> Result: The packages are added to the opt/NTAP/SANToolkit/bin directory.

Installing the Host Utilities from the CD-ROM

To install the FCP Solaris Host Utilities software from the CD-ROM, complete the following:

Step	Action
1	Login to the host system as root.
2	Mount the CD-ROM on the host system and change to the directory where the CD-ROM is mounted.
3	Install the Host Utilities package by entering the following command: <pre>pkgadd -d <attachkit_file>.pkg</pre>

Obtaining documentation

Documentation for the Host Utilities is available on the NOW site and on the FCP Solaris Host Utilities 4.2 CD-ROM. The SAN/iSAN Information Library page (<http://now.netapp.com/NOW/knowledge/docs/san/>) contains documentation for all the NetApp Host Utilities.

Completing the installation and configuring system parameters

Configuration Methods

Configuring system parameters involves modifying the `/kernel/drv/lpfc.conf` and `/etc/system` files to recommended values. You can configure system parameters by:

- ◆ Manually editing the `/kernel/drv/lpfc.conf` and `/etc/system` files
- ◆ Using the `basic_config` command. This command is provided as part of the FCP Solaris Host Utilities.

Before you begin

Complete the following pre-requisites before you configure system parameters.

- ◆ Enter the `fcv show cfmodes` command, on the storage system, to identify the storage system's `cfmode` setting. Required values in the host configuration files vary and depend on the `cfmode` setting.

If you plan to change the storage system `cfmode` to single image, first complete the procedure in Appendix C, “[Changing cfmodes to Single Image](#),” on page 81. Then follow the procedure below to configure system parameters.

- ◆ Create your own backup of the `/kernel/drv/lpfc.conf` and `/etc/system` file. The `basic_config` command automatically creates a backup for you, but does not allow you to revert more than one time.

Setting parameters manually

Procedure overview To set parameters manually you must:

- ◆ Identify the cfmode setting of the storage system
- ◆ Set /kernel/drv/lpfc.conf parameters to required values
- ◆ Set /etc/system parameters to required values

These are described in the sections that follow.

Note

For a complete list of the host parameters that NetApp recommends you change and an explanation of why the change is recommended, see the *Host Settings Affected by the Host Utilities* document, which is on the NOW site on the SAN/iSAN Information Library page (<http://now.netapp.com/NOW/knowledge/docs/san/>).

Identifying the cfmode

You must enter the `fcpl show cfmode` command, on the storage system, to identify the storage system's cfmode setting. Required values in the host configuration files vary, and depend on the cfmode.

Setting lpfc.conf variables

To set /kernel/drv/lpfc.conf variables that match the storage system cfmode, identify the cfmode and enter one of the following:

- ◆ For Standby (clustered storage systems)
 - ❖ `linkdown-tmo=30`
 - ❖ `nodev-tmo=120`
- ◆ For Single_image or Partner (clustered storage systems)
 - ❖ `linkdown-tmo=30`
 - ❖ `nodev-tmo=30`
- ◆ For Stand Alone (single storage system)
 - ❖ `linkdown-tmo=255`
 - ❖ `nodev-tmo=255`

To reset the topology parameter, enter one of the following:

- ❖ topology=2 - for switch (point to point)
- ❖ topology=4 - for arbitrated loop (this is a direct-attach configuration)

To set values for all other parameters, enter:

- ◆ tgt-queue-depth=32
- ◆ lun-queue-depth=16
- ◆ scan-down=0
- ◆ num-iocbs=1024
- ◆ num-bufs=1024
- ◆ dqfull-throttle-up-time=4
- ◆ dqfull-throttle-up-inc=2

Setting /etc/system variables

To set /etc/system variables that match the storage system cfmode, identify your cfmode and enter one of the following:

- ◆ For Standby (clustered storage systems), set the value to 120 seconds:
 - ❖ set sd:sd_io_time=0x78
- ◆ For Single_image or Partner (clustered storage systems), set the value to 30 seconds:
 - ❖ set sd:sd_io_time=0x1e
- ◆ For Stand-Alone (single storage systems), set the value to 255 seconds:
 - ❖ set sd:sd_io_time=0xff

If your host is Solaris 9 and you are configuring the FCP Solaris Host Utilities, set the following:

- ◆ set sd:sd_retry_count=5

Using the basic_config command

The basic_config command

The basic_config command is a configuration script that you can use to:

- ◆ Update items in the /kernel/drv/lpfc.conf to the values recommended for NetApp storage systems
- ◆ Identify your Solaris OS version, and set recommended values for that version in the /etc/system file
- ◆ Update timeout parameters in the /etc/system and /kernel/drv/lpfc.conf file, based on the cfmodes that you specify

The basic_config command displays the recommended values for your approval before the file is updated. It also creates backup copies of the /etc/system and /kernel/drv/lpfc.conf files.

Configuring system parameters

To configure system parameters, follow the steps below.

Step	Action
1	To display existing values, enter the following: basic_config -p Result: The software displays the current setting for NetApp required items in the /etc/system and /kernel/drv/lpfc.conf file

Step	Action
2	To configure /etc/system and /kernel/drv/lpfc.conf file items to the recommended setting, enter: basic_config -m <cfmode> where:-m describes the cfmode. Arguments are: ❖ single_image ❖ partner ❖ dual_fabric ❖ standby ❖ mixed ❖ stand_alone Result: The software displays recommended settings and prompts you for confirmation.
3	Enter y to confirm the changes and update the files. Example: This example uses the basic_config -m single_image command. The command displays the variables it will reset. Then it asks you to confirm this action. <pre> basic_config -m single_image /etc/system file will have these changes: set sd:sd_io_time=0x1e Emulex lpfc.conf will have these changes: lun-queue-depth=16; tgt-queue-depth=32; scan-down=0; num-iocbs=1024; num-bufs=1024; dqfull-throttle-up-time=4; dqfull-throttle-up-inc=2; nodev-tmo=30; linkdown-tmo=30; Please confirm the changes? (yes no) yes Please reboot your host for the changes to take effect. </pre>

Additional options The basic_config command has the following options.

Option	Description	Notes
-l	For direct attached configurations, sets the topology parameter in the /kernel/drv/lpfc.conf file to loop (topology=4).	The topology option in the /kernel/drv/lpfc.conf file is set to the proper value automatically when you configure the host. Use this option to reset the topology only if you change your configuration after you install.
-s	For switched configurations, sets the topology parameter in the /kernel/drv/lpfc.conf file to switch (topology=2).	
-r	Reverts changes to the /kernel/drv/lpfc.conf and /etc/system files made by the basic_config script. This is done by copying the backup file to the current file.	Important: The basic_config script allows you to revert only one time. To revert more than one time, you must create your own backup copy of the /kernel/drv/lpfc.conf and /etc/system file.
-v	Displays version information for the basic_config script.	

Removing an FCP Solaris Host Utilities or Attach Kit installation

Methods for removing an FCP Solaris Attach Kit

If you had installed an FCP Solaris Host Utilities or Attach Kit on the host in the past, you must remove it.

The method you use for uninstalling that installation depends on the version of the kit that you installed.

FCP Solaris Host Utilities 4.x: Use the following command to remove the FCP Solaris Host Utilities SAN Toolkit software:

```
# pkgrm NTAPSANTool
```

FCP Solaris Attach Kit 3.0: Use the following command to remove FCP Solaris Attach Kit 3.0 software from the /opt/NTAP/SANToolkit/bin directory.

```
# pkgrm NTAPSANTool
```

If you no longer need the lpfc driver and any /etc/system settings that may have been installed with basic_config, you should remove them.

FCP Solaris Attach Kit 2.0: Use the uninstall script from the FCP Solaris Attach Kit 2.0 software to uninstall previous versions of the Attach Kit software. The uninstall script automatically creates a backup copy of the /kernel/drv/lpfc.conf and sd.conf files as part of the uninstall procedure. However, it is a good practice to create a separate backup copy before you begin the uninstall.

Uninstalling the Attach Kit 2.0 software

To uninstall the Attach Kit 2.0 software, complete the following steps.

Step	Action
1	Ensure that you are logged in as root.
2	Locate the FCP Solaris Attach Kit 2.0 software. By default, the FCP Solaris Attach Kit is installed in /opt/NTAP/sanlun/bin.
3	From the /opt/NTAP/sanlun/bin directory, enter the following command to remove the existing software: ./uninstall

Step	Action
4	At the prompt “Are you sure you want to uninstall lpfc and sanlun packages?”enter y. Result: The uninstall script creates a backup copy of the /kernel/drv/lpfc.conf and sd.conf files to /usr/tmp. The backup files are named as follows: ◆ lpfc.conf.save ◆ sd.conf.save If a backup copy already exists, the install script prompts you to overwrite the backup copy.
5	Reboot the host with the reconfigure option (reboot -- -r).

About this chapter This chapter describes what you need to do to complete the configuration.

Topics in this chapter This chapter includes the following topics:

- ◆ [“Configuration overview”](#) on page 34
- ◆ [“Mapping igroups and LUNs”](#) on page 36
- ◆ [“Creating persistent bindings”](#) on page 37
- ◆ [“Configuring the host to discover LUNs”](#) on page 40
- ◆ [“Labeling the LUN and configuring Volume Management software”](#) on page 42

Configuration overview

Task overview

Completing the configuration involves the following tasks.

Task	Discussion	Additional information
1. Create and map igroups and LUNs	An igroup is a collection of WWPNs on the storage system that map to one or more host HBAs. After you create the igroup, you must create LUNs on the storage system, and map the LUNs to the igroup.	For complete information, refer to your version of the <i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml .
2. Create persistent bindings	Persistent bindings permanently bind a particular target ID on the host to the storage system WWPN. You must create persistent binding between the storage system (target) and the Solaris host (initiator) to guarantee that the storage system is always available at the correct SCSI target ID on the host.	Refer to “ Creating WWPN bindings with lputil ” on page 38 for information about creating bindings with <code>lputil</code> . Refer to “ Using the create_binding.pl Script ” on page 45 for information about creating bindings with the <code>create_binding.pl</code> script.
3. Modify the sd.conf file	The <code>sd.conf</code> file is a SCSI driver configuration file that tells the host system which disks to probe when the system is booted. You must create an entry for every NetApp LUN in the <code>sd.conf</code> file. You must also add entries to this file if you add additional LUNs.	Refer to “ Configuring the host to discover LUNs ” on page 40 for additional information.
4. Discover LUNs	LUNs are discovered when you perform a reconfiguration reboot (<code>reboot - - -r</code>) on the host. You can also discover LUNs using the <code>/usr/sbin/devfsadm</code> command.	Refer to “ Discovering LUNs ” on page 41 for additional information.

Task	Discussion	Additional information
5. Label LUNs	Use the Solaris format utility to label the LUNs.	Refer to Solaris documentation for complete information about the format utility.
6. Configure Volume Management software	When your configuration includes Veritas Volume Management software, you must use Veritas tools and utilities to bring the new LUNs under Veritas control.	Refer to Veritas documentation for additional information.

Mapping igroups and LUNs

Methods for creating igroups and LUNs

You can create igroups and LUNs on the storage system by:

- ◆ Entering the `lun setup` command
This method prompts you through the process of creating a LUN, creating an igroup, and mapping the LUN to the igroup.
- ◆ Entering a series of individual commands (such as `lun create`, `sanlun fcp show adapter -c`, `igroup create` and `lun map`).
Use this method to create one or more LUNs and igroups in any order.
- ◆ Using FilerView
This method provides a LUN Wizard that steps you through the process of creating and mapping new LUNs.

For complete information about creating igroups and LUNs, refer to *Data ONTAP™ Block Access Management Guide for FCP* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

Creating persistent bindings

What are persistent bindings?

Persistent bindings map the WWPN of each active port on the storage system to a target ID on the Solaris host initiator (HBA). Creating persistent bindings guarantees that the storage system is always available at the correct SCSI target ID on the host. Without persistent bindings, the SCSI target ID is not guaranteed.

The FCP Solaris Host Utilities supports persistent bindings between a WWPN on the storage system and a target ID on the host.

Methods for creating persistent bindings

You can use any of the following tools and methods to create persistent bindings:

- ◆ Emulex Light Pulse Utility (`lputil`) — This is an installation and configuration script that loads automatically when you install the Emulex driver. Refer to “[Creating WWPN bindings with lputil](#)” on page 38 for information about creating persistent bindings using `lputil`.
- ◆ Emulex HBAnyware — The is an HBA management GUI, provided by Emulex, and included with the driver software. To configure persistent binding with HBAnyware, refer to the documentation that came with your Emulex software.
- ◆ The `create_binding.pl` script — This is a NetApp provided script that loads when you install the FCP Solaris Host Utilities software.
 - ❖ To install the FCP Solaris Host Utilities software, refer to “[Installing the FCP Solaris Host Utilities](#)” on page 22
 - ❖ To configure persistent bindings using the `create_binding.pl` script, refer to “[Creating WWPN bindings with create_binding.pl](#)” on page 50.
- ◆ Manually editing the `kernel/drv/lpfc.conf` file.

You should create a backup of the `/kernel/drv/lpfc.conf` file before you create persistent bindings.

Creating WWPN bindings with lputil

Follow the steps below to create persistent bindings using `lputil`. Make a note of the output; you will use that information when you configure the host to discover LUNs.

Caution

Previous versions of the FCP Solaris Host Utilities may set the `automap` parameter in the `/kernel/drv/lpfc.conf` file to 0 (`automap=0`). You must reset this parameter to 1 (`automap=1`) and reboot the host before you use `lputil` to create persistent bindings.

Step	Action
1	On the Solaris host, enter the following: <pre>/usr/sbin/lputil</pre> Result : The system activates the <code>lputil</code> program, and displays the Main Menu.
2	At the Enter choice prompt, select 5 to configure persistent bindings. Result : <code>lputil</code> shows the Persistent Bindings Menu and prompts you for a binding method.
3	At the Persistent Bindings Menu, select 5 to bind automapped targets. Result: <code>lputil</code> displays a list of adapters.
4	Select the number next to the adapter for which you will configure persistent bindings. Result: <code>lputil</code> displays the automapped WWPN and WWNN associated with the adapter you specify.

Step	Action
5	Press Enter or Return to continue. Enter yes at the following prompt if the WWPN names displayed in the previous step match the WWPN names for the initiator: Bind all auto-mapped targets? y Result: lputil displays Select Binding Method options. Note _____ The FCP Solaris Host Utilities supports WWNN bindings only on filers in standby mode.
6	To bind by port name, enter 2. Result: lputil creates persistent WWPN bindings for the adapter you specify and displays the Persistent Bindings Menu.
7	To display the new bindings, enter 2 .
8	Repeat steps 3 through 7 to configure persistent bindings for additional adapters.
9	Quit lputil. 1. Enter 0 to return to the Main Menu. 2. Enter 0 to exit lputil.
10	Select one of the actions below: ◆ To continue configuration, complete the procedure in “Modifying the sd.conf file” on page 40 ◆ To quit lputil and save the changes you have made, reboot the host with the reconfigure option (reboot -- -r) Note _____ If you have not set up the sd.conf file, you set it up and then reboot your system. See “Modifying the sd.conf file” on page 40.

Configuring the host to discover LUNs

Modifying the sd.conf file

You must add an entry for each new target ID in the sd.conf file so the host system knows which disks to probe when the system is rebooted.

To modify the sd.conf file, follow the steps below.

Step	Action
1	Create a backup copy of the /kernel/drv/sd.conf file.
2	With a text editor, open the /kernel/drv/sd.conf file on the host and create an entry for the target ID mapped to the storage system's WWPN. The entry has the following format: <pre>name="sd" parent="lpfc" target=0 lun=0;</pre> You must create an entry for every LUN mapped to the host. Note_____ If the host is connected to multiple storage systems or clusters, you must create entries in sd.conf for each additional storage system or cluster, and for each LUN mapped to the host. _____
3	On the host console, reboot the host with the reconfigure option (reboot --r). Note_____ If you're performing multiple configuration steps that require rebooting, you may want to perform all those steps and then reboot. _____

Discovering LUNs

To discover new LUNs and validate that they are visible on the host, complete the following step.

Step	Action
1	To discover new LUNs, select one of the methods below: ◆ Reboot the host with the reconfigure option (<code>reboot -- -r</code>) ◆ Enter <code>/usr/sbin/devfsadm</code> Note— Occasionally the <code>/usr/sbin/devfsadm</code> command does not find LUNs. If this occurs, reboot the host with the reconfigure option. Result: The system probes for new devices. When it finds the new LUNs, it might generate a warning about a corrupt label. This warning means that the host discovered new LUNs that need to be labeled as Solaris disks.

Labeling the LUN and configuring Volume Management software

Labeling the new LUN

The format utility is a menu driven script that is provided on the Solaris host. Use the format utility to format and label new LUNs.

Follow the step below to activate the format utility and configure LUNs.

Step	Action
1	To label LUNs for the Solaris host: 1. /usr/sbin/format 2. At the format> prompt, select the disk you want to modify. 3. When prompted to label the disk, enter y. The LUN is now labeled and ready for the volume manager to use. 4. When you finish, you can use the quit option to exit the utility. Example 1: This example labels disk number 5 on a SPARC system. (Portions of this example have been removed to make it easier to review.) <pre>\$ format Searching for disks...done AVAILABLE DISK SELECTIONS: 0. c1t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@0,0 1. c1t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@1,0 2. c1t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@2,0 3. c1t3d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@3,0 4. c3t0d0 <NETAPP-LUN-0.2 cyl 6142 alt 2 hd 16 sec 640> /pci@1c,600000/lpfc@1/sd@0,0 5. c4t0d0 <NETAPP-LUN-0.2 cyl 6142 alt 2 hd 16 sec 640> /pci@1d,700000/lpfc@2/sd@0,0 Specify disk (enter its number): 5 selecting c4t0d0 [disk formatted] ... Disk not labeled. Label it now? y</pre>

Configuring Volume Management

When your configuration uses Volume Management software, you must configure the LUNs so they are under the control of the Volume Manager.

- ◆ Use the Veritas `vxctl enable` command to bring new LUNs under Veritas control
- ◆ Use the Veritas `vxdiskadm` utility to manage existing disk groups

For additional information, refer to the documentation that shipped with your Volume Management software.

Appendix overview

This appendix describes how to use the `create_binding.pl` script to create persistent bindings. This script is bundled with the FCP Solaris Host Utilities software.

You can also create bindings using `lputil` or HBAnyware.

- ◆ To create persistent bindings with `lputil`, refer to “[Creating WWPN bindings with lputil](#)” on page 38
- ◆ To create persistent bindings with HBAnyware, refer to HBAnyware documentation

Topics in this appendix

This appendix includes the following topics:

- ◆ “[The create_binding.pl script](#)” on page 46
- ◆ “[Before you begin](#)” on page 48
- ◆ “[Using the create_binding.pl script](#)” on page 50

The create_binding.pl script

What is the create_binding.pl script

The `create_binding.pl` script is a NetApp provided tool that enables you to automatically bind the WWPNs of a single storage system's active ports to a target ID and saves these bindings in the `/kernel/drv/lpfc.conf` file.

When you upgrade an installation that uses WWNN bindings to WWPN bindings, the `create_binding.pl` script saves the existing `/kernel/drv/lpfc.conf` file to `/kernel/drv/lpfc.conf.NTAP.[date-stamp]`. The script then deletes the WWNN or WWPN bindings in the `/kernel/drv/lpfc.conf` file and replaces them with WWPN bindings to the new target IDs. The previous WWNN or WWPN bindings remain in the `/kernel/drv/lpfc.conf` file but are commented out. You can refer to the commented out settings for information about your previous bindings.

For new installations in which there are no existing WWPN bindings, the `create_binding.pl` script accesses the storage system's WWPN list and automatically binds the WWPNs to target IDs and updates the `/kernel/drv/lpfc.conf` file.

Rules the script uses for WWPN bindings

The script uses the following rules for creating WWPN bindings:

- ◆ Initiators bind to accessible targets only.
- ◆ For multipath configurations with Veritas VxVM and DMP software, the script automatically creates bindings for all accessible paths to the storage system.
- ◆ On non-multipathing hosts, an initiator cannot bind to more than one active port per storage system. For these configurations, the script prompts you to select one binding for one accessible path to the storage system for each HBA.
- ◆ When the storage system's FCP cluster failover mode (cfmode) is set to standby, the script creates bindings for only the active A ports on the storage systems in a cluster. Standby mode means that the B ports are not available to the initiator until a takeover occurs.
- ◆ When the storage system's FCP cluster failover mode is set to single image, the script creates bindings for all available target ports.
- ◆ The script has two options for configurations that have multiple paths to the storage system.
 - d—enables you to configure bindings for all available paths when Veritas VxVM with DMP is installed.

-D—enables you to configure bindings for all available paths when Veritas VxVM with DMP is not yet installed.

Caution

All multipath configurations require Veritas VxVM with DMP. If you use the -D option to create bindings and you do not have Veritas software, you must install Veritas software after you create bindings. You cannot create or use LUN data until Veritas software is installed.

When the script does not automate bindings

The `create_binding.pl` script does not automatically update WWPN bindings if you already have existing WWPN bindings saved in the `/kernel/drv/lpfc.conf` file. The script instead provides information about the bindings that it finds. You might have existing bindings in the following situations:

- ◆ You have one host connected to multiple storage systems or storage clusters, and you already used the script to create bindings for one storage system or one storage cluster.
- ◆ You replaced a failed HBA on the storage system and installed a new HBA.
- ◆ You have WWPN bindings from a previous configuration.

In these situations, you can use the `create_binding.pl` script to obtain information about available paths from the host to the storage system. The script detects which storage system ports are connected to host HBA ports and displays the WWPN bindings for these paths on the host console. You can add these WWPN bindings to the `/kernel/drv/lpfc.conf` file yourself. You can edit the file manually or use the `lputil` utility to add these bindings.

Note

You should always create a backup of the `/kernel/drv/lpfc.conf` file before you edit it manually or change it using a configuration script.

Before you begin

Prerequisites for using the `create_binding.pl` script

To use the `create_binding.pl` script, your configuration must meet the following prerequisites:

- ◆ The Emulex software is installed.
- ◆ The FCP Solaris Host Utilities is installed.

The `create_binding.pl` script is bundled with the FCP Solaris Host Utilities software. You must install this kit to obtain the `create_binding.pl` software.

Refer to “[Installing the FCP Solaris Host Utilities](#)” on page 22 for information about downloading and installing the FCP Solaris Host Utilities.

- ◆ Cabling is connected

The `create_binding.pl` script checks the availability of paths from the host to the storage system. All cables must be connected correctly between the host system and the storage system.

- ◆ Zoning is configured

If zoning is configured on your switches, the zones must be configured correctly and must be active. You can verify that zones are configured correctly in the following ways:

- ❖ By checking the zoning information on the switch to verify that the storage system ports and the host ports are visible to the switch.
- ❖ By verifying that storage system ports are visible to the host. You can run the `fcplib show initiator` command on the storage system console to view the WWPNs of initiators connected to the storage system’s target HBAs.

- ◆ Storage clusters are configured

The `create_binding.pl` script configures WWPN bindings for both storage systems in a cluster configuration. If you have a storage cluster configuration, the cluster must be set up and configured before you run the script.

Storage system prerequisites

To use the `create_binding.pl` script, your storage system must meet the following prerequisites:

- ◆ Remote Shell (`rsh`) service is installed
The script uses the `rsh` service to access the storage system. You enable the `rsh` service by entering the following command:
`options rsh.enable.on`
- ◆ The FCP service is licensed and started on the storage system.
- ◆ The SAN Solaris host is a trusted host of the storage system. You use the following storage system command to make the SAN Solaris host a trusted host:
`options trusted.hosts <hostname>`

Using the create_binding.pl script

Creating WWPN bindings with create_binding.pl

To use the `create_binding.pl` script create WWPN bindings, complete the following steps.

Step	Action
1	Ensure that you are logged in to the host with root privileges.
2	Change to the <code>/opt/NTAP/SANToolkit/bin</code> directory: <code>cd /opt/NTAP/SANToolkit/bin</code>

Step	Action
3	Enter the following command: <pre>create_binding.pl -l 'login:passwd, login:passwd' -n filer -d -D</pre> -l <i>login:passwd</i> is the login name and password used to log on to the storage system. Use a quoted, comma separated pair to specify two different login:passwd strings for a storage system and its partner. -n <i>storage_system</i> is the storage system's host name. For cluster configurations, you specify one member of the cluster. The script detects and configures bindings for the partner storage system. -d enables the script to process information for Veritas DMP configurations. Use the -d option if you have already installed Veritas Volume Manager (VxVM) with DMP. If you use the -d option for DMP configurations, the script automatically configures WWPN bindings for all accessible paths to the storage system. -D enables the script to create bindings for all accessible paths to the storage system. Use this option if you have multiple paths to the storage system, but have not yet installed Veritas VxVM with DMP. Caution The -D option enables you to configure persistent bindings even if you have not yet installed Veritas software. If you have multiple paths to the storage system, you must install Veritas software after you configure persistent bindings. Multipathing software is required for configurations with multiple paths to the storage system. Example : <pre>./create_binding.pl -l root -n storage_system_1</pre>

Step	Action
4	Take one of the actions below: ◆ If you have a single-attached or direct-attached configuration, the script displays the new WWPN bindings. • Enter y to save the new WWPN bindings and proceed to Step 5 ◆ If you have VxVM with DMP and you did not use the -d option to the <code>create_binding.pl</code> script in Step 3, the script displays a warning because you did not use the -d option. The <code>create_binding.pl</code> script will not discover multiple paths to the storage system: • Enter n to exit the script • Rerun the <code>create_binding.pl</code> script with the -d option • Enter y at the prompt to continue the script and bind by WWPN Result: The script displays the new WWPN bindings and prompts you to save them to the <code>/kernel/drv/lpfc.conf</code> file. • Enter y to save the new WWPN bindings • Proceed to Step 5 ◆ If you have VxVM and DMP installed without multipathing configuration and you did not use the -d option of the script in Step 3, the script displays a warning message. • At the warning, enter y to continue. • For each HBA, choose one of the available paths to the storage system. Each path is a WWPN binding. • When prompted, enter y to save the new WWPN bindings to the <code>/kernel/drv/lpfc.conf</code> file. • Proceed to Step 5.

Step	Action
5	With a text editor, open the /kernel/drv/sd.conf file on the host and create an entry for the target ID mapped to the storage system's WWPN. The entry has the following format: <pre>name="sd" parent="lpfc" target=0 lun=0;</pre> You must create an entry for every LUN mapped to the host. Note_____ If the host is connected to multiple storage systems or clusters, you must create entries in sd.conf for each additional storage system or cluster, and for each LUN mapped to the host. _____
6	On the host console, reboot the host with the reconfigure option (reboot --r). Note_____ If you're performing multiple configuration steps that require rebooting, you may want to perform all those steps and then reboot. _____

About this appendix This appendix describes what you need to do to configure SAN boot on a Solaris host.

Note

NetApp qualifies solutions and components on an ongoing basis. To verify that SAN booting is supported in your configuration, see the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products*, which is online at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml.

Topics in this appendix

Topics in this appendix include the following:

- ◆ “[SAN boot overview](#)” on page 56
- ◆ “[Configuration overview](#)” on page 57
- ◆ “[Setting up the Emulex HBA](#)” on page 59
- ◆ “[Copy bootblks and data](#)” on page 72
- ◆ “[Modifying OpenBoot](#)” on page 75

SAN boot overview

What is SAN boot

SAN boot is using a SAN attached disk, like a LUN, as a boot device for a Solaris host. Configuring SAN boot allows you to do the following:

- ◆ Remove the hard drives from your servers and use the SAN for your boot needs, thereby eliminating the cost associated with maintaining and servicing hard drives
- ◆ Use the SAN for both storage and boot needs, thereby consolidating and centralizing storage

Contents of this appendix

This appendix provides a procedural example that uses a Solaris host to configure a SAN boot LUN and install a new OS image onto it. The procedure described in this appendix is an example of one method for creating a bootable LUN. You can apply this procedure to most configurations.

This appendix does not describe other examples of bootable LUNs, such as configurations that boot multiple volumes or use diskless servers.

Refer to Solaris documentation for details about additional configuration methods.

Sun documentation for this procedure

For more information on the procedures that follow, see Sun's document *Sun StorEdge SAN Foundation Software Configuration Guide*, which is available in PDF format on the Sun site.

Configuration overview

Prerequisites

The procedure and example in this chapter assume the following:

- ◆ The FCP Solaris Host Utilities software has been installed and the host and storage system are configured properly and use software and firmware that is supported by the Host Utilities.
- ◆ The host operating system is installed on a local disk and uses a UFS file system.
- ◆ The OpenBoot firmware (Bootcode) is downloaded and installed on the HBA.

The OpenBoot firmware is available on the Emulex site. Follow steps in “[Upgrading firmware](#)” on page 18 for details on how to load this firmware onto the HBA.
- ◆ The Emulex FCA utilities with EMLXemlxu and the EMLXdrv have been installed. For more information on these tools, see “[Install the EMLXemlxu utilities](#)” on page 14.

For information about host, storage system and firmware versions, refer to the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

Configuration overview

Configuring a bootable LUN involves:

1. Setting up the HBA
 - ❖ Use the emlxdrv tool to remove the EMLXS driver binding with the HBA
 - ❖ Shut down the system and set the Emulex HBA to use SD mode
 - ❖ Use the emlxdrv tool to bind the Emulex HBA with the lpfc driver
2. Creating the boot LUN
 - ❖ Create the LUN that you will use for the bootable LUN.
 - ❖ Display the size and layout of the partitions of the current Solaris boot drive
 - ❖ Partition the bootable LUN to match the host boot drive
3. Copying the bootblks and boot data
 - ❖ Install the bootblks onto the bootable LUN
 - ❖ Copy the boot data from the source disk onto the bootable LUN

4. Modifying OpenBoot
 - ❖ Verify the OpenBoot version on the adapter
 - ❖ Set the FCP topology to the bootable LUN
 - ❖ Bind the adapter target and the bootable LUN
 - ❖ Create an alias for the bootable LUN
5. Rebooting the system.

Setting up the Emulex HBA

Setting up the Emulex HBA

Part of configuring a bootable LUN involves setting up your Emulex HBA. You need to perform the following tasks:

- ◆ Use the `emlxdrv` tool to unbind the Emulex HBA from the EMLXS driver. For Solaris 10 update 1 and later, the `emlxs` driver that comes with the Sun StorEdge SAN Foundation Software automatically binds to the Emulex HBA. You must use the `clear_emlxs` command, which comes with the Emulex `emlxdrv` tool, to remove these bindings.
- ◆ Shut down the system and set the Emulex HBA to use SD mode
- ◆ Use the `emlxdrv` tool to bind the Emulex HBA to the `lpfc` driver

Unbinding the HBA Complete the following steps to unbind the Emulex HBA from the emlxs driver.

Step	Action
1	Start the <code>emlxdrv</code> tool, which is in <code>/opt/EMLXemlxu/bin/emlxdrv</code>. When you start this tool, it displays information about the drivers, which are using emlxs names. Note _____ By default, when the SFS is enabled, it gives the emlxs names to the Emulex HBA devices. Example: Executing the <code>emlxdrv</code> tool checks the device names and displays a list of commands available with the tool. <pre> #/opt/EMLXemlxu/bin/emlxdrv EMLXDRV Driver Management Utility, Version 1.00j COPYRIGHT © 2004-2005 Emulex. All rights reserved. Driver Alias Present Boot Sun Models ----- emlxs lpfs no no no LP8000S and LP9002S (SBUS) emlxs f800 no no no LP8000 and LP8000DC emlxs f900 no no no LP9002, LP9002C, LP9002DC, and LP9402DC emlxs f980 no no no LP9802 and LP9802DC emlxs fa00 yes no no LP10000, LP10000DC and LP10000ExDC emlxs fd00 no no no LPe11000 and LPe11002 emlxs fe00 no no no LPe11000 and LPe11002 emlxs f0a5 no no no 2G Blade Adapter (emlxs only) emlxs fc00 no no yes LP10000-S and LP10000DC-S emlxs fc20 no no yes LPe11000-S and LPe11002-S Available commands: ----- set_emlxs <Alias> - Sets emlxs driver to bind to the specified device(s) set_emlxs_sun - Sets emlxs driver to bind to all Sun devices set_emlxs_all - Sets emlxs driver to bind to all devices set_lpfc <Alias> - Sets lpfc driver to bind to the specified device(s) set_lpfc_nonsun - Sets lpfc driver to bind to all non-Sun devices clear_dev <Alias> - Clears driver binding to the specified device(s) clear_lpfc - Clears all lpfc driver bindings clear_emlxs - Clears all emlxs driver bindings clear_sun - Clears driver bindings to all Sun devices clear_nonsun - Clears driver bindings to all non-Sun devices clear_all - Clears driver bindings to all devices q - Exits this program. </pre>

Step	Action
2	Enter the <code>clear_emlxs</code> command to remove the emlxs driver bindings from the HBA. Example: Executing the <code>clear_emlxs</code> command removes the bindings. As you can see from the output of the command, the Emulux HBA no longer uses emlxs names. <pre># emlxdv> clear_emlxs Driver Alias Present Boot Sun Models ----- - lpfs no no no LP8000S and LP9002S (SBUS) - f800 no no no LP8000 and LP8000DC - f900 no no no LP9002, LP9002C, LP9002DC, and LP9402DC - f980 no no no LP9802 and LP9802DC - fa00 yes no no LP10000, LP10000DC and LP10000ExDC - fd00 no no no LP11000 and LP11002 - fe00 no no no LPe11000 and LPe11002 - f0a5 no no no 2G Blade Adapter (emlxs only) - fc00 no no yes LP10000-S and LP10000DC-S - fc20 no no yes LPe11000-S and LPe11002-S</pre>
3	Use the <code>set_lpfc_nonsun</code> command to set the HBA to an lpfc driver. Example: You start the emlxdv tool from the <code>/opt/EMLXemlxu/bin/</code> directory. <pre>#emlxdv> set_lpfc_nonsun</pre>
4	Enter <code>q</code> to quit the emlxdv tool.

Changing the HBA to sd mode

To change the HBA from SFS compatible mode to SD mode, you must bring the system down and then change each HBA.

Complete the following steps.

Step	Action
1	At the operating system prompt issue the <code>init 0</code> command.
2	At the ok prompt, enter the <code>reset-all</code> command.

Step	Action
3	Issue the <code>show-devs</code> command to see the current device names. Example: Executing the <code>show-devs</code> shows that the Emulex device has not been set to SD mode. If the mode had been changed, the devices would have displayed <code>.../lpfc</code> instead of <code>.../emlx</code>. See Step 4 for information setting the devices to SD mode. <pre> ok show-devs controller@1,400000 /SUNW,UltraSPARC-III+@1,0 /memory-controller@0,400000 /SUNW,UltraSPARC-III+@0,0 /virtual-memory /memory@m0,0 /aliases /options /openprom /chosen /packages /pci@8,600000/SUNW,qlc@4 /pci@8,600000/SUNW,qlc@4/fp@0,0 /pci@8,600000/SUNW,qlc@4/fp@0,0/disk /pci@8,700000/emlx@3 /pci@8,700000/emlx@1 /pci@8,700000/scsi@6,1 /pci@8,700000/scsi@6 /pci@8,700000/usb@5,3 </pre>

Step	Action
4	Select the first Emulex device and set it to SD mode using the set-sd-boot command. This changes the Emulex devices to lpfc devices. Example: The first command, <code>select</code>, selects the Emulex device 700000/emlx@1. The second command, <code>set -sd-boot</code>, sets the SD mode. <pre> ok select /pci@8,700000/emlx@1 ok set-sd-boot Flash data structure updated. Signature 4e45504f Valid_flag 4a Host_did 0 Enable_flag 5 SFS_Support 0 Topology_flag 0 Link_Speed_flag 0 Diag_Switch 0 Boot_id 0 Lnk_timer f Plogi-timer 0 LUN (1 byte) 0 DID 0 WWPN 0000.0000.0000.0000 LUN (8 bytes) 0000.0000.0000.0000 *** Type reset-all to update. *** </pre>
5	Repeat the previous step for each Emulex device.
6	Enter the reset-all command to update the devices. Example: The <code>reset</code> command update the Emulex devices with the new mode. <pre> ok reset-all Resetting ... </pre>

Step	Action
7	Issue the show-devs command to confirm that you have changed the mode on all the Emulex devices. Example: Executing the show-devs command lets you confirm that all the Emulex devices are showing up as lpfc devices. <pre> ok show-devs /pci@8,600000 /pci@8,700000 /memory-controller@1,400000 /SUNW,UltraSPARC-III+@1,0 /memory-controller@0,400000 /SUNW,UltraSPARC-III+@0,0 /virtual-memory /memory@m0,0 /aliases /options /openprom /chosen /packages /pci@8,600000/SUNW,qlc@4 /pci@8,600000/SUNW,qlc@4/fp@0,0 /pci@8,600000/SUNW,qlc@4/fp@0,0/disk /pci@8,700000/lpfc@3 /pci@8,700000/lpfc@1 /pci@8,700000/scsi@6,1 /pci@8,700000/scsi@6 /pci@8,700000/usb@5,3 </pre>
8	Use the boot command to bring the system back up. Example: The boot command brings the system back up. <pre> ok boot Boot device: /pci@8,600000/SUNW,qlc@4/fp@0,0/disk@w21000004cfafe0cf,0:a File an d args: SunOS Release 5.10 Version Generic_118833-33 64-bit Copyright 1983-2006 Sun Microsystems, Inc. All rights reserved. Use is subject to license terms. \ </pre>

Step	Action
9	Use the <code>sanlun</code> command to check the Emulex HBA device names to confirm that they show up as <code>lpfc</code> devices. Example: The <code>sanlun</code> command displays information about the Emulex HBA device names. <pre># sanlun fcp show adapter -v</pre> <pre> adapter name: lpfc1 WWPN: 10000000c94bfb30 WWNN: 20000000c94bfb30 driver name: lpfc model: LPe11002-M4 model description: Emulex LPe11002-M4 4Gb 2port FC: PCIe SFF HBA serial number: VM54033886 hardware version: 2057706d driver version: 6.02h; HBAAPI(I) v2.0.e, 11-07-03 firmware version: 2.50A6 (Z2F2.50A6) Number of ports: 1 port type: Fabric port state: Operational supported speed: 4 GBit/sec negotiated speed: 4 GBit/sec OS device name: /devices/pci@1e,600000/pci@0/pci@8/lpfc@0 adapter name: lpfc2 WWPN: 10000000c94bfb31 WWNN: 20000000c94bfb31 driver name: lpfc model: LPe11002-M4 model description: Emulex LPe11002-M4 4Gb 2port FC: PCIe SFF HBA serial number: VM54033886 hardware version: 2057706d driver version: 6.02h; HBAAPI(I) v2.0.e, 11-07-03 firmware version: 2.50A6 (Z2F2.50A6) Number of ports: 1 port type: Fabric port state: Operational supported speed: 4 GBit/sec negotiated speed: 4 GBit/sec OS device name: /devices/pci@1e,600000/pci@0/pci@8/lpfc@0,1 </pre>

Creating the bootable LUN

Create the LUN

Use standard storage system commands and procedures to create and map a LUN assigned to Host ID 0. The example in this appendix uses host side LUN 0, (c5t3d0s2) as the bootable LUN.

For detailed information about configuring igroups and LUNs, refer to *Data ONTAP® Software Setup Guide* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

Partition the bootable LUN to match the host device

You must partition the bootable LUN so that it matches the partitions on the host boot device. This involves:

- ◆ Displaying information about the host boot device
- ◆ Modifying the bootable LUN to model the partition layout of the host boot device

Complete the following steps to partition the bootable LUN.

Step	Action
1	Select the current Solaris boot device: 1. Activate the Solaris format utility. Enter /usr/sbin/format. Result: The system displays Available Disk Selections. Example: The format utility displays information about the disks on your system. <pre>#format Searching for disks...done AVAILABLE DISK SELECTIONS: 0. c1t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfafe0cf,0 1. c1t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfafe220,0 2. c5t0d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@1/sd@0,0 3. c5t1d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@1/sd@1,0 4. c5t2d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@1/sd@2,0 5. c5t3d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@1/sd@3,0 6. c6t0d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@3/sd@0,0 7. c6t1d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@3/sd@1,0 8. c6t2d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@3/sd@2,0 9. c6t3d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@3/sd@3,0</pre> Specify disk (enter its number): Note— If you want to set up additional SAN boot configurations, you can use the <code>san lun</code> command to find the primary path. The examples in this section use the <code>c5t3d0s2</code> device. 2. At the specify disk prompt, enter the number that corresponds to the current Solaris boot device. This is the device that maps to the root file system in the <code>/etc/vfstab</code> file. Result: The system selects the Solaris boot disk and displays partition options.

Step	Action																																																						
2	Display partition information for the current Solaris boot disk: 1. At the <code>format></code> prompt, enter p to display partition menu options 2. At the <code>partition></code> prompt, enter p to print the partition table for the current Solaris boot device. Result: The system displays the partitions for the boot device. Example: The following is the partition information for the current root disk slice being used for the SAN booting example in this chapter. Current partition table (original): Total disk cylinders available: 14087 + 2 (reserved cylinders) <table><tr><th>Part</th><th>Tag</th><th>Flag</th><th>Cylinders</th><th>Size</th><th>Blocks</th></tr><tr><td>0</td><td>root</td><td>wm</td><td>207 - 2267</td><td>10.00GB</td><td>(2061/0/0) 20972736</td></tr><tr><td>1</td><td>swap</td><td>wu</td><td>0 - 206</td><td>1.00GB</td><td>(207/0/0) 2106432</td></tr><tr><td>2</td><td>backup</td><td>wm</td><td>0 - 14086</td><td>68.35GB</td><td>(14087/0/0) 143349312</td></tr><tr><td>3</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>4</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>5</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>6</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>7</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr></table>	Part	Tag	Flag	Cylinders	Size	Blocks	0	root	wm	207 - 2267	10.00GB	(2061/0/0) 20972736	1	swap	wu	0 - 206	1.00GB	(207/0/0) 2106432	2	backup	wm	0 - 14086	68.35GB	(14087/0/0) 143349312	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	unassigned	wm	0	0	(0/0/0) 0	7	unassigned	wm	0	0	(0/0/0) 0
Part	Tag	Flag	Cylinders	Size	Blocks																																																		
0	root	wm	207 - 2267	10.00GB	(2061/0/0) 20972736																																																		
1	swap	wu	0 - 206	1.00GB	(207/0/0) 2106432																																																		
2	backup	wm	0 - 14086	68.35GB	(14087/0/0) 143349312																																																		
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	unassigned	wm	0	0	(0/0/0) 0																																																		
7	unassigned	wm	0	0	(0/0/0) 0																																																		
3	Determine the size and layout of the partitions of the source boot volume. Make a note of the following: ◆ Partition size - the bootable LUN must be the same size or larger than the original boot volume. If your system requires Veritas, you must allow an additional 1 GB for Veritas encapsulation. ◆ Partition layout -the bootable LUN must contain the same partitions in the same locations as the host boot device(s). Devices you must include are referenced in the <code>vfstab</code> file. Note_____ If your configuration boots off of more than one device, repeat steps 1 through 3 to determine the size and layout of each device. You must create and configure a bootable LUN that matches each boot device on the host. _____																																																						

Step	Action																																																						
	Example: The following shows the contents of the /etc/ vfstab file. Devices c5t3d0s1 and c5t3d0s0 (in bold) reference the source boot device and must be replicated on the bootable LUN. <pre>#vi /mnt/etc/vfstab #device device mount FS fsck mount mount #to mount to fsck point type pass at boot options # fd - /dev/fd fd - no - /proc - /proc proc - no - /dev/dsk/c5t3d0s1 - - swap - no - /dev/dsk/c5t3d0s0 /dev/rdisk/c5t3d0s0 / ufs 1 no - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes -</pre>																																																						
4	Use format to create the same size root and swap slice for the SAN boot disk. Example: This example shows the output of the print option for the current boot device on the Solaris host. It uses 10 GB for root and 1 GB for swap.Partitions that correspond to the vfstab file are displayed in bold. Current partition table (unnamed): Total disk cylinders available: 6398 + 2 (reserved cylinders) <table><tr><th>Part</th><th>Tag</th><th>Flag</th><th>Cylinders</th><th>Size</th><th>Blocks</th></tr><tr><td>0</td><td>root</td><td>wm</td><td>32 - 351</td><td>10.00GB</td><td>(320/0/0) 20971520</td></tr><tr><td>1</td><td>swap</td><td>wu</td><td>0 - 31</td><td>1024.00MB</td><td>(32/0/0) 2097152</td></tr><tr><td>2</td><td>backup</td><td>wm</td><td>0 - 6397</td><td>199.94GB</td><td>(6398/0/0) 419299328</td></tr><tr><td>3</td><td>unassigned</td><td>wu</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>4</td><td>unassigned</td><td>wu</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>5</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>6</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>7</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr></table>	Part	Tag	Flag	Cylinders	Size	Blocks	0	root	wm	32 - 351	10.00GB	(320/0/0) 20971520	1	swap	wu	0 - 31	1024.00MB	(32/0/0) 2097152	2	backup	wm	0 - 6397	199.94GB	(6398/0/0) 419299328	3	unassigned	wu	0	0	(0/0/0) 0	4	unassigned	wu	0	0	(0/0/0) 0	5	unassigned	wm	0	0	(0/0/0) 0	6	unassigned	wm	0	0	(0/0/0) 0	7	unassigned	wm	0	0	(0/0/0) 0
Part	Tag	Flag	Cylinders	Size	Blocks																																																		
0	root	wm	32 - 351	10.00GB	(320/0/0) 20971520																																																		
1	swap	wu	0 - 31	1024.00MB	(32/0/0) 2097152																																																		
2	backup	wm	0 - 6397	199.94GB	(6398/0/0) 419299328																																																		
3	unassigned	wu	0	0	(0/0/0) 0																																																		
4	unassigned	wu	0	0	(0/0/0) 0																																																		
5	unassigned	wm	0	0	(0/0/0) 0																																																		
6	unassigned	wm	0	0	(0/0/0) 0																																																		
7	unassigned	wm	0	0	(0/0/0) 0																																																		
5	Once you have created the slices, use the format command to label the disks.																																																						

Step	Action
6	Select the bootable LUN: 1. Navigate to the Available Disk Selections screen: ❖ At the <code>partition></code> prompt enter q ❖ At the <code>format></code> prompt enter q Result: The system displays the Available Disk Selections screen. 2. At the <code>specify disk</code> prompt, enter the number that corresponds to the LUN mapped to host controller and Target 0. Note_____ Make a note of the identifying information for the LUN you select. You will use this information in “Modifying OpenBoot” on page 75, when you bind the bootable LUN to the adapter and create a boot alias. _____ Example: The following example selects c5t3d0 as the bootable LUN. <pre>#format Searching for disks...done AVAILABLE DISK SELECTIONS: 0. c1t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfafe0cf,0 1. c1t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w21000004cfafe220,0 2. c5t0d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@1/sd@0,0 3. c5t1d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@1/sd@1,0 4. c5t2d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@1/sd@2,0 5. c5t3d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@1/sd@3,0 6. c6t0d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@3/sd@0,0 7. c6t1d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@3/sd@1,0 8. c6t2d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@3/sd@2,0 9. c6t3d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 4096> /pci@8,700000/lpfc@3/sd@3,0 Specify disk (enter its number): Specify disk (enter its number): 5 selecting c5t3d0s2 [disk formatted]</pre>

Step	Action
7	To display information about the bootable LUN: 1. At the format> prompt, enter p to display partition menu options. 2. At the partition> prompt, enter p to print the partition table for the current Solaris boot device. Result: The system displays the partitions for the bootable LUN
8	At the partition> prompt, enter modify. Modify the partitions on the bootable LUN to match the host boot device.
9	Exit the format utility. 1. At the partition> prompt, enter q. 2. At the format> prompt, enter q.

Copy bootblks and data

- Install the bootblk

The bootblk contains startup information, required by the Solaris host, that you must install onto the raw bootable LUN. Installing the bootblk involves:

 - ◆ Using the `uname -a` command to determine the directory in `/usr/platform` where the bootblk is located
 - ◆ Installing the bootblk onto the bootable LUN.

Complete the following steps to install the bootblk onto the bootable LUN.

Step	Action
1	Determine the directory in <code>/usr/platform</code> where the bootblks are located. Enter: uname -a Example: The output from the example that follows indicates that bootblks should reside in <code>/usr/platform/sun4u</code> (bold). <pre>#uname -a SunOS shsun17 5.10 Generic_118833-33 sun4u sparc SUNW,Sun-Fire-280R</pre>
2	Write the bootblk. Enter: <pre>/usr/sbin/installboot /usr/platform/<directory_name>/lib/fs/ufs/bootblk /dev/rdisk/<device_name></pre> Example: The following example installs bootblks in the <code>/usr/platform/sun4u</code> directory on <code>/dev/rdsk/c5t3d0s0</code>. <pre>#/usr/sbin/installboot /usr/platform/sun4u/lib/fs/ufs/bootblk /dev/rdsk/c5t3d0s0</pre>

- Copy the boot data

Copying the boot data to the bootable LUN involves

 - ◆ Creating and mounting a file system on the bootable LUN. The file system must match the file system type on the host boot device. The example in this document uses a UFS filesystem.
 - ◆ Moving boot data from the host device onto the bootable LUN
 - ◆ Editing the `vfstab` file to reference the bootable LUN

Complete the following steps to copy the boot data onto the bootable LUN.

Step	Action
1	Create a file system on the bootable LUN that matches the file system in the vfstab file of the source device. Enter: <pre>newfs <device></pre> Example: The following creates a USF filesystem at /dev/rdisk/c5t3d0s0. <pre>#newfs /dev/rdisk/c5t3d0s0 newfs: /dev/rdisk/c5t3d0s0 last mounted as / newfs: construct a new file system /dev/rdisk/c5t3d0s0: (y/n)? y Warning: 4096 sector(s) in last cylinder unallocated /dev/rdisk/c5t3d0s0: 20971520 sectors in 3414 cylinders of 48 tracks, 128 sectors 10240.0MB in 214 cyl groups (16 c/g, 48.00MB/g, 5824 i/g) super-block backups (for fsck -F ufs -o b=#) at: 32, 98464, 196896, 295328, 393760, 492192, 590624, 689056, 787488, 885920, 20055584, 20154016, 20252448, 20350880, 20449312, 20547744, 20646176, 20744608, 20843040, 20941472</pre> The next example shows the corresponding entry (in bold) for the file system in the vfstab file. <pre>#device device mount FS fsck mount mount #to mount to fsck pont type pass at boot options fd - /dev/fd fd - no - /proc - /proc proc - no - /dev/dsk/c5t3d0s1 - - swap - no - /dev/dsk/c5t3d0s0 /dev/rdisk/c5t3d0s0 / ufs 1 no - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes -</pre>
2	Mount the root directory of the bootable LUN. Enter: <pre>mount <filesystem_name> /<mountpoint></pre> Example: The following mounts the file system on the bootable LUN under /mnt. <pre>#mount /dev/dsk/c5t3d0s0 /mnt</pre>

Step	Action
3	Create the required directory structure onto the bootable LUN and copy the boot data. You need to perform a dump and a restore from the current boot disk to the SAN disk. You can use the commands: <pre>ufsdump 0f - <source_boot_device> cd/ <mntpoint of bootable_lun>; ufsrestore rf -)</pre> Example: The following brings the system down to single user mode and performs a dump of the current boot disk. It then restore it to the SAN boot disk. <pre>#init S #df -k Filesystem kbytes used avail capacity Mounted on /dev/dsk/c1t0d0s0 10327132 3882023 6341838 38% / #ufsdump 0f - /dev/dsk/c1t0d0s0 (cd /mnt; ufsrestore rf -) DUMP: Date of this level 0 dump: Wed 28 Mar 2007 02:22:29 PM EDT DUMP: Date of last level 0 dump: the epoch DUMP: Dumping /dev/rdisk/c1t0d0s0 (shsun17:/) to standard output. DUMP: Mapping (Pass I) [regular files] DUMP: Mapping (Pass II) [directories] DUMP: Writing 32 Kilobyte records DUMP: Estimated 8004176 blocks (3908.29MB). DUMP: Dumping (Pass III) [directories] DUMP: Dumping (Pass IV) [regular files] Warning: ./lost+found: File exists DUMP: 98.97% done, finished in 0:00 DUMP: 8004158 blocks (3908.28MB) on 1 volume at 6506 KB/sec DUMP: DUMP IS DONE</pre>
4	Edit the vfstab file to refer to the bootable LUN. Example: The following shows the vfstab entry for the bootable LUN (bold). <pre>#vi /mnt/etc/vfstab #more /mnt/etc/vfstab #device device mount FS fsck mount mount #to mount to fsck point type pass at boot options # fd - /dev/fd fd - no - /proc - /proc proc - no - /dev/dsk/c5t3d0s1 - - swap - no - /dev/dsk/c5t3d0s0 /dev/rdisk/c5t3d0s0 / ufs 1 no - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes -</pre>

Modifying OpenBoot

What is OpenBoot

OpenBoot is the firmware that the Emulex software uses to startup the system. OpenBoot firmware also includes the hardware level user interface that you use to configure the bootable LUN.

Modifying OpenBoot

Modifying OpenBoot involves:

- ◆ Validating the OpenBoot version
- ◆ Setting the topology (point-to-point or arbitrated loop) to the bootable LUN
- ◆ Binding the HBA to the bootable LUN
- ◆ Creating an alias for the bootable LUN. This alias substitutes for the device address (/pci@8,700000/lpfc@1/sd@3,0 in the example that follows) during subsequent boot operations.

Complete the following to modify OpenBoot.

Step	Action
1	Activate the OpenBoot environment. Enter: halt - to halt the Solaris Operating system Result: The OpenBoot program displays the OpenBoot prompt. setenv auto-boot? false - to stop automatic reboot during OpenBoot configuration reset-all - reset the host and activate your changes
2	Check the OpenBoot version that is currently installed on the host adapter. Enter: select /<lpfc_device> - to select the host adapter .version - to display the OpenBoot version For information about the correct firmware version, refer to the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/. unselect-dev - to unselect the host adapter

Step	Action
3	Set the topology to the bootable LUN. - At the OpenBoot prompt, enter probe-scsi-all Result: The system displays all scsi devices - Select the adapter associated with the bootable LUN. Enter: ❖ select <boot_LUN_pathname> - Set the network topology. Enter one of the following: ❖ set-ptp - to set point to point ❖ set-fc-al - to set arbitrated loop. - Enter reset-all to reset the host and activate your changes.

Step	Action
	Example: The following example displays SCSI information. It uses selects lpfc@1 as the adapter that accesses the bootable LUN 0 on WWPN 500a098483e0b481. The topology is set to point-to-point. <pre> {1} ok probe-scsi-all /pci@6,2000/lpfc@1 Device PortID 610e13 WWPN 500a098483e0b481 LUN 0 Disk NETAPP LUN 0.2 LUN 1 Disk NETAPP LUN 0.2 LUN 2 Disk NETAPP LUN 0.2 LUN 3 Disk NETAPP LUN 0.2 Device PortID 610f13 WWPN 500a098383e0b481 LUN 0 Disk NETAPP LUN 0.2 LUN 1 Disk NETAPP LUN 0.2 LUN 2 Disk NETAPP LUN 0.2 LUN 3 Disk NETAPP LUN 0.2 /pci@1f,4000/scsi@2 Target 6 Unit 0 Removable Read Only device TOSHIBA DVD-ROM SD-M14011007 /pci@1f,4000/scsi@3 Target 0 Unit 0 Disk SEAGATE ST336704LSUN36G 032C Target 2 Unit 0 Disk IBM DDYST1835SUN18G S94N {0} ok select /pci@6,2000/lpfc@1 {0} ok .nvram Signature 4e45504f Valid_flag 4a Host_did 0 Enable_flag 5 SFS_Support 0 Topology_flag 0 Link_Speed_flag 0 Diag_Switch 0 Boot_id 0 Lnk_timer f Plogi-timer 0 LUN (1 byte) 0 DID 0 WWPN 0000.0000.0000.0000 LUN (8 bytes) 0000.0000.0000.0000 </pre>

Step	Action
4	Bind the host adapter to the bootable LUN. 1. Select the adapter associated with the bootable LUN. Enter: ❖ select <boot_LUN_pathname> 2. Bind the bootable LUN wwpn, target, and LUN using the set-boot-id command.

Step	Action
	Example: The following example binds the host adapter lpfc@1 to LUN 0. <pre> {1} ok probe-scsi-all /pci@6,2000/lpfc@1 Device PortID 610e13 WWPN 500a098483e0b481 LUN 0 Disk NETAPP LUN 0.2 LUN 1 Disk NETAPP LUN 0.2 LUN 2 Disk NETAPP LUN 0.2 LUN 3 Disk NETAPP LUN 0.2 Device PortID 610f13 WWPN 500a098383e0b481 LUN 0 Disk NETAPP LUN 0.2 LUN 1 Disk NETAPP LUN 0.2 LUN 2 Disk NETAPP LUN 0.2 LUN 3 Disk NETAPP LUN 0.2 /pci@1f,4000/scsi@2 Target 6 Unit 0 Removable Read Only device TOSHIBA DVD-ROM SD-M14011007 /pci@1f,4000/scsi@3 Target 0 Unit 0 Disk SEAGATE ST336704LSUN36G 032C Target 2 Unit 0 Disk IBM DDYST1835SUN18G S94N {1} ok select /pci@6,2000/lpfc@1 {{0} ok WWPN 500a098483e0b481 0 3 set-boot-id Flash data structure updated. Signature 4e45504f Valid_flag 14e Host_did 0 Enable_flag 5 SFS_Support 0 Topology_flag 0 Link_Speed_flag 0 Diag_Switch 0 Boot_id 3 Lnk_timer f Plogi-timer 0 LUN (1 byte) 0 DID 0 WWPN WWPN 500a098483e0b481 LUN (8 bytes) 0000.0000.0000.0000 </pre>

Step	Action
5	Create an alias for the bootable LUN. Enter: <pre>nvalias <alias_name> /<boot_LUN_pathname></pre> Example: This example creates an alias called “sanbootdisk”. <pre>{0} ok nvalias sanbootdisk /pci@8,700000/lpfc@1/sd@3,0 {0} ok nvstore {0} ok devalias sanbootdisk /pci@8,700000/lpfc@1/sd@3,0</pre>
6	To allow the system to autoboot to the LUN, add the LUN’s alias to your boot device list. Enter the following command: <pre>setenv boot-device <alias_name></pre> Note_____ Before adding the new alias, use the <code>printenv boot-device</code> command to see if there are already boot devices in the list. If aliases are in the list, you’ll need to include them on the <code>setenv boot-device</code> command line after the bootable LUN’s alias. That alias must always be the first argument. The format for the command would be: <pre>setenv boot-device <bootable_lun_alias_name> <existing_alias1> <existing_alias2> ...</pre> _____ Example: You’ve created the alias <code>sanbootdisk</code> for your bootable LUN. You use the check the boot device list and discover that there are already two boot devices in the list: <code>disk</code> and <code>net</code>. To add the alias <code>sanbootdisk</code>, you would enter the <code>setenv boot-device</code> command followed by the <code>sanbootdisk</code> and the other two aliases in the list. <pre>{0} ok setenv boot-device sanbootdisk disk net</pre>
7	Boot the LUN. Enter: <pre>boot <alias_name></pre> Example: The following command boots the alias for <code>/pci@1c,600000/lpfc@1/sd@0,0</code>. <pre>{0} ok boot sanbootdisk</pre>

Appendix overview This appendix describes how to upgrade your host to single image mode.

Topics in this appendix This appendix includes the following topics:

- ◆ “[Understanding cfmode](#)” on page 82
- ◆ “[Completing prerequisites](#)” on page 84
- ◆ “[Upgrading to single image mode](#)” on page 85

Understanding cfmode

What is single image mode?

Single image is a cfmode setting that is available as of Data ONTAP 7.1 and later. In single image mode, both systems in the cluster share the same World Wide Node Name (WWNN).

Single image mode provides the following advantages:

The host can access all LUNs through any target port on the cluster.

- ◆ Single image mode is supported on all NetApp storage systems and by all supported FCP hosts.
- ◆ There are no switch limitations. You can connect an NetApp storage system in single image mode to any supported FCP switch.

Reasons for using single image mode

You might want to migrate your configuration to single image mode for the following reasons:

- ◆ To increase the number of paths to the cluster
You can increase the number of available paths in your configuration because all ports on each storage system are available to the host.
- ◆ To support heterogeneous configurations
For example, you upgraded an existing SAN configuration with Solaris hosts and you want to add HP-UX or AIX hosts to the configuration.

Host requirements

HP-UX and AIX hosts do not support failover when the storage systems in a cluster are in standby mode. If you want to add these hosts to an existing homogenous configuration that has Solaris-only hosts, you have to change the cfmode setting on the storage system to either single image or partner.

Solaris hosts in the same environment as HP-UX or AIX hosts must also have multipathing software installed when the cfmode setting is single image or partner.

Impact of changing the cfmode setting

Changing the cfmode setting of a cluster impacts the following components of your configuration:

- ◆ Access to LUNs
When you change the cfmode setting, you change the available paths between the host and the storage cluster. Some previously available paths

no longer available and some new paths become available. You must reconfigure every host that is connected to the cluster to discover the new paths. The LUNs might be accessible but must not be used until you reconfigure the host.

- ◆ Multipathing policy

If you have multipathing software in your configuration, changing the `cfmode` setting might also affect the multipathing policy because there are more paths between the host and the cluster.

- ◆ Zoning

Updating the filer to single image `cfmode` changes the WWNN of the cluster and the World Wide Port Names (WWPNs) of the target ports on each storage system. If the cluster connects to a switch by using soft zoning (zoning by WWPN), you must check the zoning setup and make the required updates to your switch.

- ◆ Switch configuration and cabling

In mixed and dual fabric modes, the target ports log into the fabric in loop mode. In single image mode, the target ports log into the fabric using point-to-point (PTP) mode. This means you might need to update your switch configuration when you migrate from these modes to single image mode so that the target connects to the fabric as point-to-point device. For additional information, see the *Data ONTAP® Block Access Management Guide for FCP* at

http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

Completing prerequisites

Before you begin

You should complete the following prerequisite tasks before you change your cfmode.

Task	What to do
1. Check system requirements	You must verify that your system uses the supported version of driver and volume management software. Refer to the latest configuration information, on the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml.
2. Schedule the downtime	Hosts cannot access LUNs while you change cfmode. You must allow sufficient downtime to change the storage system cfmode, and reconfigure and reboot the host.
3. Save configuration information	You must record and save system configuration information before you begin. On the host, make a backup of all the volume groups that contain LUNs. These backups must contain up-to-date information for both the data within the volume group and the volume group configuration. Enter and save output from the following commands: ◆ <code>san lun lun show</code> - display LUNs visible to the host. ◆ <code>san lun lun show -p</code> - display LUNs and paths visible to the host If your LUNs are under Veritas disk management, save the output from the following commands: ◆ <code>vxdisk list</code> - display disks under Veritas control ◆ <code>vxdisk path</code> - display disks and paths under Veritas control ◆ <code>vxprint -htg diskgroup_name</code> - display Veritas volumes and status
4. Identify affected hosts and initiators	You must reconfigure and reboot any hosts and initiators that access LUNs on the storage systems you are upgrading. To identify affected hosts and initiators, enter the following command: ◆ <code>igroup show</code>
5. Take the host applications off-line	Enter commands that quiesce host I/O and take host applications off-line.

Upgrading to single image mode

Procedure overview

To upgrade an FCP attached Solaris host and storage system to single image mode you must complete the following tasks:

- 1. Prepare the host —This involves stopping I/O and protecting the configuration by unmounting file systems and deporting Veritas disk groups.
- 2. Upgrade the storage system to single image mode — This involves setting the cfmode to single image. You may also need to resolve LUN mapping or OS-type conflicts, if they occur.
- 3. Validate zoning — This involves checking zoning and resolving zone mapping issues.
- 4. Reconfigure the host — This involves generating persistent bindings and updating the sd.conf file. You must also reconfigure parameters in the /kernel/drv/lpfc.conf and /etc/system files to the values required by the storage system.
- 5. Reboot the host — This involves a reconfiguration reboot of the system. When your configuration uses VxVM volume management, you must also enter additional commands to access Veritas disk groups and file systems.

These are described in the sections that follow.

Preparing the host

To prepare the host for upgrade to single image mode, complete the following steps.

Step	Action
1	Open a window to the Solaris host
2	Enter commands that quiesce I/O and halt applications on each storage system that you want to upgrade.

Step	Action
3	If you are using Veritas VxVM software, unmount the file systems and deport the associated Veritas disk groups. Enter: 1. <code>/usr/sbin/umount /mountpoint</code> — to unmount the file system 2. <code>/usr/sbin/vxdg deport diskgroup_name</code> — to deport the Veritas diskgroup
4	Verify that the Veritas diskgroup is deported. Enter: <code>/usr/sbin/vxdisk -o alldgs list</code> Result: Veritas displays exported disk groups in parentheses.

Upgrading the storage cluster to single image mode

To upgrade the storage cluster to single image mode, complete the following steps.

Step	Action
1	Open a window to the console of each storage system you want to upgrade.
2	Login at the console of the storage system(s) you are upgrading, and enter the following command to identify and resolve LUN and igroup mapping conflicts. <code>lun config_check -S</code> Result: Mappings that are displayed using the “*” character indicate an error, and also provide LUN and igroup information that you can use to resolve conflicts. Use the <code>lun map</code> command to re-map LUNS to a new ID that is unique to each node in the cluster. You use the <code>igroup</code> command to change igroup otypes, if required.
3	Switch to advanced privileges by entering the following command: <code>priv set advanced</code>
4	Stop the FCP service. On each storage system, enter: <code>fcp stop</code>

Step	Action
5	Set the cfmode to single_image. On each storage system, enter: <pre>fcp set cfmode single_image</pre> If you see messages indicating mapping conflicts, then return to Step 2 to resolve the mapping conflicts.
6	Start the FCP service on each system in the cluster. On each system, enter the following command: <pre>fcp start</pre>

Validating zoning and connectivity

Changes to WWNNs and WWPNS: After you change the cfmode setting to single image both systems in the cluster share a single WWNN. One node in the cluster takes the WWNN of its partner. This means that the WWNN of one of the systems in the cluster has changed. The WWPNS of the target ports also change. You might have to update your zoning configuration depending on the type of zoning you use. Soft zoning uses the WWPNS. Port zoning uses the physical location (switch port) in the zone. If you use soft zoning you must update your zones to use the new WWPNS. Configurations that use hard zoning are not impacted.

Cabling and zoning: If you did not change your cabling, migrating to single image mode also changes the number of available paths to the LUNs. Some additional paths to the target become available because more target ports are visible to the host. You can use zoning or Data ONTAP portsets to reduce the number of paths in your configuration to comply with the number of paths supported by your multipathing software.

For details about validating and modifying switch zones see the documentation provided by your switch manufacturer.

For details about Data ONTAP portsets, see the [Data ONTAP Block Access Management Guide](#).

You might also want to change your cabling if you have more paths between the host and the cluster and want to reduce the number of paths. It is recommended that after you change the cfmode setting, make sure each HBA has one primary and at least one secondary path to each LUN.

Reconfiguring the host

To reconfigure the host, complete the following steps.

Step	Action
1	Because changing the cfmode to single image changes the WWPN of the target HBA's, you must update the persistent bindings in the /kernel/drv/lpfc.conf file. You can use opt/NTAP/SANToolkit/bin/create_bindings.pl or usr/sbin/lpfc/lputil to re-generate persistent bindings in the /kernel/drv/lpfc.conf file. Select one of the options below. ◆ To use lputil enter: lputil From the lputil Main Menu, select 5 Follow screen prompts or refer to “Creating persistent bindings” on page 37 for additional information about using lputil. ◆ To use create_bindings.pl, enter: create_binding.pl -d -n storage_system_name Refer to Appendix A, “Using the create_binding.pl Script,” on page 45 for additional information about using create_binding.pl.
2	With a text editor, open the /kernel/drv/sd.conf file on the host and create an entry for the target ID mapped to the storage system's WWPN. The entry has the following format: <pre>name="sd" parent="lpfc" target=0 lun=0;</pre> You must create an entry for every LUN mapped to the host. If the host is connected to multiple storage systems or clusters, you must create entries in sd.conf for each additional storage system or cluster, and for each LUN mapped to the host. Note _____ You should create a backup of the sd.conf file before you make changes. _____
3	Update /etc/system and /kernel/drv/lpfc.conf parameters to the recommended values. Enter: basic_config -m single_image

Completing the upgrade

Follow the steps below complete the cfmode upgrade.

Step	Action
1	Perform a reconfigure reboot (reboot -- -r) on the host.
2	When your configuration uses a Volume Manager, you must import the Veritas diskgroup and mount each file system. Enter: ◆ /usr/sbin/vxdg import <i>diskgroup_name</i> - to import the Veritas diskgroup ◆ /usr/sbin/mount /mountpoint - to mount the file system

Appendix overview

If you are using Veritas Storage Foundation for multipathing, you must install and configure the Symantec Array Support Library (ASL) for NetApp storage systems. This chapter describes how to install and configure the ASL and how this component affects multipathing and failover.

Note

To determine which version of the Veritas Storage Foundation and ASL to use, check the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

Topics in this appendix

Topics in this appendix include the following:

- ◆ “[Overview of the Array Support Library](#)” on page 92
- ◆ “[Installing the ASL](#)” on page 93
- ◆ “[Removing the ASL](#)” on page 96
- ◆ “[Obtaining multipathing information](#)” on page 97

Overview of the Array Support Library

What the Array Support Library is

The Array Support Library (ASL) is an ONTAP-qualified library that provides information about storage array attributes configuration to the Device Discovery Layer (DDL) and Mucomponent of VxVM.

The DDL is a component of VxVM that discovers available enclosure information for disks and disk arrays that are connected to a host system. The DDL calls ASL functions during the storage discovery process on the host. The ASL in turn “claims” a device based on vendor and product identifiers. The claim associates the storage array model and product identifiers with the device.

You can add or remove ASLs from a running VxVM system. You do not need to reboot the host.

Note

In a Storage Foundation RAC cluster, you must stop clustering on a node before you remove the ASL.

What the ASL provides

The ASL provides enclosure-based naming information about SAN-attached storage systems. With enclosure-based naming, the name of the disk is based on the logical name of its enclosure, or disk array, and not a raw device name. The ASL provides specific information to VxVM about SAN-attached storage systems, instead of referring to them as JBOD devices or raw devices.

By using the ASL, you obtain the following information about the LUNs:

- ◆ Enclosure name—VxVM's enclosure-based naming feature creates disk names based on the name of its enclosure, or disk array, and not a raw device name.
- ◆ Multipathing policy—Whether the NetApp storage is accessed as an Active/Active (A/A) disk array or an Active/Passive Group Concurrent (A/PG-C) disk array. The ASL also provides information about primary and secondary paths to NetApp storage.

Using storage with the ASL

You cannot use storage systems simultaneously as JBODs and vendor arrays. If you install the ASL, storage systems cannot be configured in VxVM as JBODs. They are reported as storage arrays unless you explicitly exclude them by using the `vxddladm exclude array` command.

Installing the ASL

ASL installation overview

Installing the ASL involves the following tasks:

- ◆ Verifying that your configuration meets system requirements. See the *Compatibility and Configuration Guide for NetApp FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml> for current information about the system requirements.
- ◆ Determining the ASL version if it is currently installed.
- ◆ Obtaining the ASL.
- ◆ Installing a fresh version of the ASL or upgrading an existing ASL version.

Determining the ASL version

If you currently have the ASL installed, use the Veritas `vxddladm listversion` command to determine its version.

Example: The following shows output from the `vxddladm listversion` command.

vxddladm listversion		
LIB_NAME	ASL_VERSION	Min. VXVM version
=====		
libvxCLARiiON.so	vm-5.0-rev-1	5.0
libvxcscovrts.so	vm-5.0-rev-1	5.0
libvxemc.so	vm-5.0-rev-2	5.0
libvxengenio.so	vm-5.0-rev-1	5.0
libvxhds9980.so	vm-5.0-rev-1	5.0
libvxhdsalua.so	vm-5.0-rev-1	5.0
libvxhdsusp.so	vm-5.0-rev-2	5.0
libvxhpalua.so	vm-5.0-rev-1	5.0
libvxibmds4k.so	vm-5.0-rev-1	5.0
libvxibmds6k.so	vm-5.0-rev-1	5.0
libvxibmds8k.so	vm-5.0-rev-1	5.0
libvxsenas.so	vm-5.0-rev-1	5.0
libvxshark.so	vm-5.0-rev-1	5.0
libvxsunse3k.so	vm-5.0-rev-1	5.0
libvxsunset4.so	vm-5.0-rev-1	5.0
libvxvpath.so	vm-5.0-rev-1	5.0
libvxxp1281024.so	vm-5.0-rev-1	5.0
libvxxp12k.so	vm-5.0-rev-2	5.0
libvxibmsvc.so	vm-5.0-rev-1	5.0
libvxnetapp.so	vm-5.0-rev-0	5.0
#		

Obtaining the ASL software

The ASL is not distributed with the Host Utilities. You obtain the ASL from the Symantec Web site. Download this software before you start the installation.

For a direct link to the ASL on the Symantec Web site, see the FCP Host Compatibility Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml.

Installing ASL software you downloaded

To install a fresh version of the ASL that you downloaded from Symantec, complete the following steps.

Step	Action
1	Log in to the VxVM system as the root user.
2	If you already have NetApp storage configured as JBOD in your VxVM configuration, remove the JBOD support for NetApp by entering: vxddladm rmjbod vid=NETAPP
3	Install the ASL according to the directions provided by Symantec. You should have your LUNs set up before you install the ASL. For information on setting up LUNs, see “ Completing the Configuration ” on page 33.
4	If your host is connected to NetApp storage, you can verify the installation by locating the NetApp Enclosure Type in the output of the vxddmpadm listenclosure all command. The output shows the model name of the storage device if you are using enclosure-based naming with VxVM. vxddmpadm listenclosure all Note_____ After you verify the installation, you can use the vxddctl enable command to restart the vxconfigd daemon. You only use this command when your host is connected to NetApp storage systems. _____ Example: The vxddmpadm listenclosure all command shows the Enclosure Types as FAS3020 in this example. <pre># vxddmpadm listenclosure all ENCLR_NAME ENCLR_TYPE ENCLR_SNO STATUS ARRAY_TYPE ===== Disk Disk DISKS CONNECTED Disk FAS30200 FAS3020 1071853 CONNECTED A/PG-C</pre>

Step	Action
5	If your host is not connected to storage, use the following command: <pre>vxddladm listsupport all</pre> Example: The following is a sample of the output you see when you enter the vxddladm listsupport all command. <pre># vxddladm listsupport all LIBNAME VID ===== libvxCLARiiON.so DGC libvxcscovrts.so CSCOVRTS libvxemc.so EMC libvxengenio.so SUN libvxhds9980.so HITACHI libvxhdsalua.so HITACHI libvxhdsusp.so HITACHI libvxhpalua.so HP, COMPAQ libvxibmds4k.so IBM libvxibmds6k.so IBM libvxibmds8k.so IBM libvxsena.so SENA libvxshark.so IBM libvxsunse3k.so SUN libvxsunset4.so SUN libvxvpath.so IBM libvxxp1281024.so HP libvxxp12k.so HP libvxibmsvc.so IBM libvxnetapp.so NETAPP</pre>

Removing the ASL

Before you begin You do not need to stop any volumes created on unaffected disk arrays, such as disk arrays from other vendors, before removing ASL support. This is also true for arrays or disks in the OTHER_DISKS category. OTHER_DISKS are local non-FCP-attached disks. Volumes created on these arrays remain accessible because they are not multipathed.

Only one version of an ASL that supports Data ONTAP storage system can be installed on the host at any given time.

Removing the ASL To remove the ASL, complete the following steps.

Step	Action
1	Log into the host as the root user.
2	Stop all I/O to NetApp LUNs. Note _____ In a Storage Foundation RAC cluster, you must also stop clustering on a node before you remove the ASL. _____
3	Enter the following command to remove the ASL package: <code>pkgrm <asl_filename></code> Example: The following command removes a previous version of the ASL. <code># pkgrm NTAPasl</code>

Obtaining multipathing information

What an ASL array type is

The ASL reports information about the multipathing configuration to the DDL as one of the following disk array types:

- ◆ Active/Active (A/A)—There are multiple active paths to storage and simultaneous I/O is supported on all paths. If a path fails, I/O is distributed across the remaining paths.
- ◆ Active/Passive Group Concurrent (A/PG-C) - An A/PG-C array is a variant of the A/PG array type that supports concurrent I/O and load balancing by having multiple primary paths to LUNs. Failover to the secondary (passive) path occurs only if all the active primary paths fail.

For additional information, see the following documents:

- ◆ *Veritas Array Support Library Developer's Guide* - for details about load balancing.
- ◆ *Veritas Volume Manager Administrator's Guide* - for details about system management

How the storage system's FCP failover mode corresponds to array types

In clustered NetApp storage configurations, the array type corresponds to the storage system cfmode settings. If you use the standby cfmode, the array type will be A/A; otherwise, it will be A/PG-C.

Note

The ASL also supports direct attached, non-clustered configurations, including NearStore models. These configurations have no cfmode settings. ASL reports these configurations as Active/Active (A/A) array types. For more details about array types, see the *Veritas Array Support Library Developer's Guide*.

How sanlun displays the array type

The Host Utilities SAN Toolkit includes the `sanlun` utility in the `opt/NTAP/SANToolkit/bin` directory. You use the `sanlun` utility to display information about paths to LUNs on the storage system.

- ◆ When the ASL is installed and the LUN is controlled by VxVM, the output of the `sanlun` command displays the `Multipath_Policy` as A/PG-C or A/A.
- ◆ If the LUN is not controlled by a volume manager, then the `Multipathing_Policy` is none and the `Multipathing-provider` is none.

Using VxVM to display available paths

To obtain information about available paths to a LUN using VxVM, complete the following steps

Step	Action
1	Enter the following command to view all devices: <pre>vxdisk list</pre> Result: The VxVM management interface displays the vxdisk device, type, disk, group and status. It also shows which disks are managed by VxVM. Example: The following is an example of the type of output you see with the vxdisk list command. <pre># vxdisk list DEVICE TYPE DISK GROUP STATUS Disk_0 auto:none - - online invalid Disk_1 auto:none - - online invalid FAS30201_0 auto:cdsdisk c2t4d60s2 dg2 online shared FAS30201_1 auto:cdsdisk c2t4d26s2 dg1 online shared FAS30201_2 auto:cdsdisk - - online FAS30201_3 auto:cdsdisk c2t4d61s2 dg2 online shared FAS30201_4 auto:cdsdisk c2t4d84s2 dg2 online shared</pre> (The output has been truncated to make the document easier to read.)
2	On the host console, enter the following command to display path information for the device you want: <pre>vxddmpadm getsubpaths dmpnodename=<device></pre> where <device> is the name listed under the output of the <code>vxdisk list</code> command. Example: The following is the type of output you should see. <pre># vxddmpadm getsubpaths dmpnodename=FAS30201_0 NAME STATE[A] PATH-TYPE[M] CTLR-NAME ENCLR-TYPE ENCLR-NAME ATTRS ===== = c2t4d60s2 ENABLED(A) PRIMARY c2 FAS3020 FAS30201 - c2t5d60s2 ENABLED(A) PRIMARY c2 FAS3020 FAS30201 - c2t6d60s2 ENABLED SECONDARY c2 FAS3020 FAS30201 - c2t7d60s2 ENABLED SECONDARY c2 FAS3020 FAS30201 - c3t4d60s2 ENABLED(A) PRIMARY c3 FAS3020 FAS30201 - c3t5d60s2 ENABLED(A) PRIMARY c3 FAS3020 FAS30201 - c3t6d60s2 ENABLED SECONDARY c3 FAS3020 FAS30201 - c3t7d60s2 ENABLED SECONDARY c3 FAS3020 FAS30201 -</pre>

Step	Action																																																																																											
3	Enter the following command to obtain path information for a host HBA: <pre>vxdmpadm getsubpaths ctlr=<controller_name></pre> controller_name is the controller displayed under “CTRL - NAME” in the output of the vxdmpadm getsubpaths dmpnodename command you entered in Step 2. Result: The output displays information about the paths to the filer storage (whether the path is a primary or secondary path). The output also lists the filer that the device is mapped to. Example: The following is the type of output you should see. <pre># vxdmpadm getsubpaths ctlr=c2</pre> <table><tr><th>NAME</th><th>STATE[A]</th><th>PATH-TYPE[M]</th><th>DMPNODENAME</th><th>ENCLR-TYPE</th><th>ENCLR-NAME</th><th>ATTRS</th></tr><tr><td colspan="7">=====</td></tr><tr><td colspan="7">=</td></tr><tr><td>c2t4d0s2</td><td>ENABLED(A)</td><td>PRIMARY</td><td>FAS30201_81</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d1s2</td><td>ENABLED</td><td>SECONDARY</td><td>FAS30201_80</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d2s2</td><td>ENABLED(A)</td><td>PRIMARY</td><td>FAS30201_76</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d3s2</td><td>ENABLED</td><td>SECONDARY</td><td>FAS30201_44</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d4s2</td><td>ENABLED(A)</td><td>PRIMARY</td><td>FAS30201_45</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d5s2</td><td>ENABLED</td><td>SECONDARY</td><td>FAS30201_2</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d6s2</td><td>ENABLED(A)</td><td>PRIMARY</td><td>FAS30201_82</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d7s2</td><td>ENABLED</td><td>SECONDARY</td><td>FAS30201_79</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d8s2</td><td>ENABLED(A)</td><td>PRIMARY</td><td>FAS30201_43</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d9s2</td><td>ENABLED</td><td>SECONDARY</td><td>FAS30201_41</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr></table> (The output has been truncated to make the document easier to read.)	NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS	=====							=							c2t4d0s2	ENABLED(A)	PRIMARY	FAS30201_81	FAS3020	FAS30201	-	c2t4d1s2	ENABLED	SECONDARY	FAS30201_80	FAS3020	FAS30201	-	c2t4d2s2	ENABLED(A)	PRIMARY	FAS30201_76	FAS3020	FAS30201	-	c2t4d3s2	ENABLED	SECONDARY	FAS30201_44	FAS3020	FAS30201	-	c2t4d4s2	ENABLED(A)	PRIMARY	FAS30201_45	FAS3020	FAS30201	-	c2t4d5s2	ENABLED	SECONDARY	FAS30201_2	FAS3020	FAS30201	-	c2t4d6s2	ENABLED(A)	PRIMARY	FAS30201_82	FAS3020	FAS30201	-	c2t4d7s2	ENABLED	SECONDARY	FAS30201_79	FAS3020	FAS30201	-	c2t4d8s2	ENABLED(A)	PRIMARY	FAS30201_43	FAS3020	FAS30201	-	c2t4d9s2	ENABLED	SECONDARY	FAS30201_41	FAS3020	FAS30201	-
NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS																																																																																						
=====																																																																																												
=																																																																																												
c2t4d0s2	ENABLED(A)	PRIMARY	FAS30201_81	FAS3020	FAS30201	-																																																																																						
c2t4d1s2	ENABLED	SECONDARY	FAS30201_80	FAS3020	FAS30201	-																																																																																						
c2t4d2s2	ENABLED(A)	PRIMARY	FAS30201_76	FAS3020	FAS30201	-																																																																																						
c2t4d3s2	ENABLED	SECONDARY	FAS30201_44	FAS3020	FAS30201	-																																																																																						
c2t4d4s2	ENABLED(A)	PRIMARY	FAS30201_45	FAS3020	FAS30201	-																																																																																						
c2t4d5s2	ENABLED	SECONDARY	FAS30201_2	FAS3020	FAS30201	-																																																																																						
c2t4d6s2	ENABLED(A)	PRIMARY	FAS30201_82	FAS3020	FAS30201	-																																																																																						
c2t4d7s2	ENABLED	SECONDARY	FAS30201_79	FAS3020	FAS30201	-																																																																																						
c2t4d8s2	ENABLED(A)	PRIMARY	FAS30201_43	FAS3020	FAS30201	-																																																																																						
c2t4d9s2	ENABLED	SECONDARY	FAS30201_41	FAS3020	FAS30201	-																																																																																						

Using sanlun to obtain multipathing information

To obtain multipathing information by using the san lun command, complete the following step.

Step	Action																																																															
1	On the host command line, enter the following command: <pre>sanlun lun show -p all</pre> Result: san lun displays path information for each LUN. For LUNs managed by VxVM, the output displays the Veritas multipathing policy (A/A or A/PG-C array types) and the path I/O policy. Example: The following partial output shows eight paths to one LUN managed by VxVM. <pre># sanlun lun show -p all filerX:/vol/vol1/data_lun60 (LUN 60) /dev/rdsk/c2t4d60s2 (FAS30201_0) 9g (9663676416) lun state: GOOD</pre> <table><tr><th>path state</th><th>path type</th><th>device filename</th><th>host HBA</th><th>target port</th><th>partner port</th><th>Multipath policy</th></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c2t4d60s2</td><td>lpfc0</td><td>1a</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c2t5d60s2</td><td>lpfc0</td><td>1b</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c3t5d60s2</td><td>lpfc1</td><td>4a</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c3t4d60s2</td><td>lpfc1</td><td>4b</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c2t7d60s2</td><td>lpfc0</td><td></td><td>1b</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c2t6d60s2</td><td>lpfc0</td><td></td><td>1a</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c3t6d60s2</td><td>lpfc1</td><td></td><td>4b</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c3t7d60s2</td><td>lpfc1</td><td></td><td>4a</td><td>A/PG-C</td></tr></table>	path state	path type	device filename	host HBA	target port	partner port	Multipath policy	enabled	primary	/dev/rdsk/c2t4d60s2	lpfc0	1a		A/PG-C	enabled	primary	/dev/rdsk/c2t5d60s2	lpfc0	1b		A/PG-C	enabled	primary	/dev/rdsk/c3t5d60s2	lpfc1	4a		A/PG-C	enabled	primary	/dev/rdsk/c3t4d60s2	lpfc1	4b		A/PG-C	enabled	secondary	/dev/rdsk/c2t7d60s2	lpfc0		1b	A/PG-C	enabled	secondary	/dev/rdsk/c2t6d60s2	lpfc0		1a	A/PG-C	enabled	secondary	/dev/rdsk/c3t6d60s2	lpfc1		4b	A/PG-C	enabled	secondary	/dev/rdsk/c3t7d60s2	lpfc1		4a	A/PG-C
path state	path type	device filename	host HBA	target port	partner port	Multipath policy																																																										
enabled	primary	/dev/rdsk/c2t4d60s2	lpfc0	1a		A/PG-C																																																										
enabled	primary	/dev/rdsk/c2t5d60s2	lpfc0	1b		A/PG-C																																																										
enabled	primary	/dev/rdsk/c3t5d60s2	lpfc1	4a		A/PG-C																																																										
enabled	primary	/dev/rdsk/c3t4d60s2	lpfc1	4b		A/PG-C																																																										
enabled	secondary	/dev/rdsk/c2t7d60s2	lpfc0		1b	A/PG-C																																																										
enabled	secondary	/dev/rdsk/c2t6d60s2	lpfc0		1a	A/PG-C																																																										
enabled	secondary	/dev/rdsk/c3t6d60s2	lpfc1		4b	A/PG-C																																																										
enabled	secondary	/dev/rdsk/c3t7d60s2	lpfc1		4a	A/PG-C																																																										

Optimizing the Veritas DMP Restore Daemon

The Veritas DMP restore daemon requirements

It is recommended that you set recommended values for the restore policy of the Veritas DMP restore daemon and also specify a polling interval. These settings determine the frequency that the Veritas daemon checks paths between the host and the storage system.

See the General Operational Considerations page of the FCP Solaris Host Utilities Release Notes for detailed information about recommended values. The Release Notes are online in the SAN/iSAN Information Library at <http://now.netapp.com/NOW/knowledge/docs/san/>.

Understanding ASL error message types

Understanding error message severity The ASL works silently and seamlessly with the VxVM DDL. If an error, misconfiguration, or malfunction occurs, messages from the library are logged to the console using the host’s logging facility. The following table lists the importance and severity of these messages.

Message severity	Definition	Action required
Error	Indicates that an ERROR status is being returned from the ASL to the VxVM DDL that prevents the device (LUN) from being used. The device might still appear in vxdisk list, but it is not usable.	Call Symantec Technical Support for help.
Warning	Indicates that an UNCLAIMED status is being returned. Unless claimed by a subsequent ASL, dynamic multipathing is disabled. No error is being returned but the device (LUN) might not function as expected.	
Info	Indicates that a CLAIMED status is being returned. The device functions fully, with DMP enabled, but the results seen by the user might be other than what is expected. For example, the enclosure name might change.	

This appendix describes how to use the `sanlun` utility to display information about the LUNs, paths to the LUNs, and the Solaris HBA.

Topics in this appendix

Topics in this appendix include the following:

- ◆ [“Displaying host LUN information”](#) on page 104
- ◆ [“Displaying information about the storage system paths”](#) on page 106
- ◆ [“Displaying host HBA information”](#) on page 108

Displaying host LUN information

Displaying host LUN information To display host LUN information, complete the following steps.

Step	Action
1	Ensure that you are logged in as root on the host.
2	Change to the /opt/NTAP/SANToolkit/bin directory: <code>cd /opt/NTAP/SANToolkit/bin</code>

Step	Action
3	Enter the following command to display LUN information: <pre>./sanlun lun show [-v] [-d <i>host device filename</i> <i>all</i> <i>storage system name</i> <i>storage system name:storage system pathname</i>] -p</pre> -v produces verbose output. -d is the device option and can be one of the following: ◆ <i>host device filename</i> is the special device file name for the disk on Solaris (which might represent a storage system LUN). ◆ <i>all</i> lists all storage system LUNs under /dev/rdisk. ◆ <i>storage system name</i> lists all storage system LUNs under /dev/rdisk on a specified storage system. ◆ <i>storage system name: storage system pathname</i> lists all storage system LUNs under /dev/rdisk that are connected to the storage system path name LUN on a specified storage system. -p provides information about the primary and secondary paths available to the LUN when you are using multipathing. Result: The requested LUN information is displayed. Example: The following is an example of the output you see when you use the sanlun command. <pre># ./sanlun lun show -v -d /dev/rdisk/c2t1d2s2 filer: lun-pathname device filename adapter lun size lun state filerX: /vol/vol1/hostA_lun2 /dev/rdisk/c2t1d2s2 lpfc1 150.0g (161061273088) GOOD Serial number: 2eFclndDjKlM Filer FCP nodename:50a9800055CC66DD Filer FCP portname:50a9800355CC66DD Filer adapter name: 5a Filer IP address: 10.22.110.36 Filer volume name:vol1 FSID:0xfffffc139 Filer qtree name:/vol/vol1 ID:0x0 Filer snapshot name: ID:0x0 LUN partition table permits multiprotocol access: no why: there can be only one data partition LUN has valid label: yes #</pre>

Displaying information about the storage system paths

Displaying path information

To display information about paths to the storage system, complete the following steps.

Step	Action																																			
1	Ensure that you are logged in as root on the host.																																			
2	Change to the /opt/NTAP/SANToolkit/bin directory: cd /opt/NTAP/SANToolkit/bin																																			
3	At the host command line, enter the following command to display LUN information: ./sanlun lun show -p all -v -p provides information about the primary and secondary paths available to the LUN when you are using multipathing. all lists all storage system LUNs under /dev/rdsk. -v turns on verbose mode. Example: The following example uses the -p option to display information about the paths. <pre># ./sanlun lun show -p all filer1:/vol/solaris/lun0 (LUN 0) /dev/rdsk/c3t0d0s2 (c3t0d0s2) 300g (322122547200) lun state: GOOD</pre> <table><tr><th>path state</th><th>path type</th><th>device filename</th><th>host HBA</th><th>target port</th><th>partner port</th><th>Multipath policy</th></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c3t0d0s2</td><td>lpfc0</td><td>9a</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c4t0d0s2</td><td>lpfc1</td><td>11a</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c3t2d0s2</td><td>lpfc0</td><td></td><td>9b</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c4t2d0s2</td><td>lpfc1</td><td></td><td>11b</td><td>A/PG-C</td></tr></table>	path state	path type	device filename	host HBA	target port	partner port	Multipath policy	enabled	primary	/dev/rdsk/c3t0d0s2	lpfc0	9a		A/PG-C	enabled	primary	/dev/rdsk/c4t0d0s2	lpfc1	11a		A/PG-C	enabled	secondary	/dev/rdsk/c3t2d0s2	lpfc0		9b	A/PG-C	enabled	secondary	/dev/rdsk/c4t2d0s2	lpfc1		11b	A/PG-C
path state	path type	device filename	host HBA	target port	partner port	Multipath policy																														
enabled	primary	/dev/rdsk/c3t0d0s2	lpfc0	9a		A/PG-C																														
enabled	primary	/dev/rdsk/c4t0d0s2	lpfc1	11a		A/PG-C																														
enabled	secondary	/dev/rdsk/c3t2d0s2	lpfc0		9b	A/PG-C																														
enabled	secondary	/dev/rdsk/c4t2d0s2	lpfc1		11b	A/PG-C																														

Interpreting sanlun lun show -p output

The `sanlun lun show -p` command provides the following information:

- ◆ path state—Whether the path is enabled or disabled.
- ◆ path type—Whether the path to the LUN is primary, secondary, or standby.
 - ❖ Primary paths communicate directly using the adapter on the local storage system.
 - ❖ Secondary paths are proxied to the partner storage system over the cluster interconnect.
 - ❖ Standby occurs if the `cfmode` is either mixed or standby and the path is being serviced by a partner storage system in takeover mode. Note that this case occurs only when Veritas assumes the array policy is Active/Active. If the array policy is Active/Passive and the path is being served by the partner file in the takeover mode, that path state displays as secondary.
- ◆ device filename—The special device file name for the disk on Solaris that represents the LUN.
- ◆ host HBA—The name of the initiator HBA on the host.
- ◆ local storage system port—The port that provides direct access to a LUN. This port appears as the `Xa` for `F8xx`, `gF8xx`, `FAS9xx`, `gFAS9xx`, `FAS3000`, or `FAS6000` storage systems and `X_C` for `FAS270` storage systems, where `X` is the slot number of the HBA. This is always a primary path.
- ◆ partner storage system port—The port that provides passive path failover. This port appears as `Xb` for `F8xx` or `FAS9xx` storage systems and `X_C` for `FAS270` on storage systems, where `X` is the slot number of the HBA. This is always a secondary path.

After the failover of a storage system cluster, the `sanlun lun show -p` command reports secondary paths as secondary but enabled, because these are now the active paths.

- ◆ multipathing policy—The multipathing policy is one of the following:
 - ❖ A/A (Active/Active)—DMP uses more than one path concurrently for traffic. The A/A policy indicates that the B ports on the storage system HBAs are in standby mode and become active only during a takeover. This means that there are two active paths to the cluster at any given time.
 - ❖ A/PG-C (Active/Passive Group Concurrent)—Groups of LUNs connected to a single controller failover as a single failover entity. Failover occurs at the controller level and not at the LUN level.
 - ❖ The primary and secondary controllers are each connected to a separate group of LUNS. If a single LUN in the primary controller's LUN group fails, all LUN's in that group fail over to the secondary controller's passive LUN group. This array type allows load balancing of I/O across multiple primary paths. Refer to Veritas documentation for details about load balancing.

Displaying host HBA information

Displaying host HBA information To manage the host HBA information, complete the following steps.

Step	Action
1	Ensure that you are logged in as root on the host.
2	Change to the /opt/NTAP/SANToolkit/bin directory: <code>cd /opt/NTAP/SANToolkit/bin</code>

Step	Action
3	At the host command line, enter the following command to display host HBA information: <pre>./sanlun fcp show adapter[-c] [-v] [adapter name all]</pre> -C option produces configuration instructions. -v option produces verbose output. all lists information for all FCP adapters. adapter name lists information for a specified adapter. Result: The FCP adapter information is displayed. Example: The following command line displays information about the adapter. <pre>#./sanlun fcp show adapter -v lpfc0 WWPN:10000000c94ae06e WWNN: 20000000c94ae06e driver name: lpfc model: LP11002-M4 model description: Emulex LP11002 4Gb/s 2-Port FC : PCI-X Mode-2 HBA SFF Diag Optics serial number: BG53164700 hardware version: 1036406d driver version: 6.02h; HBAAPI(I) v2.0.e, 11-07-03 firmware version: 2.10A7 (B2F2.10A7) Number of ports: 1 port type: Fabric port state: Operational supported speed: Unknown negotiated speed: 2 GBit/sec OS device name: /devices/pci@1d,700000/fibre-channel@1 lpfc1 WWPN:10000000c94ae06d WWNN: 20000000c94ae06d driver name: lpfc model: LP11002-M4 model description: Emulex LP11002 4Gb/s 2-Port FC : PCI-X Mode-2 HBA SFF Diag Optics serial number: BG53164700 hardware version: 1036406d driver version: 6.02h; HBAAPI(I) v2.0.e, 11-07-03 firmware version: 2.10A7 (B2F2.10A7) Number of ports: 1 port type: Fabric port state: Operational supported speed: Unknown negotiated speed: 2 GBit/sec OS device name: /devices/pci@1d,700000/fibre-channel@1,1</pre>

About this appendix This appendix provides tips and troubleshooting information to help you as you use the FCP Solaris Host Utilities. This appendix contains the following:

- ◆ Suggestions for checking your system.
- ◆ Information about possible problems and how to handle them.
- ◆ Diagnostic tools that you can use to gather information about your system.

Topics in this appendix

This chapter discusses the following topics:

- ◆ [“Confirming that your system is correctly set up”](#) on page 112
- ◆ [“Using Diagnostic Utilities”](#) on page 118
- ◆ [“Collecting storage system information”](#) on page 119
- ◆ [“Collecting Solaris information”](#) on page 121
- ◆ [“Collecting switch information”](#) on page 122

Confirming that your system is correctly set up

About the sections that follow

If you have any problems with the Host Utilities, you should make sure your system setup is correct.

The information in the sections that follow helps you check your system setup. As you go through it, keep the following in mind:

- ◆ As an experienced Solaris user, you may already know these commands. This is just a quick refresher of some of the key areas to check to avoid having problems.
For more information on any of the Solaris commands, see the man pages and operating system documentation.
For more information on the Data ONTAP commands, see the Data ONTAP documentation, in particular, the FCP and iSCSI Block Access Guide.
- ◆ You perform some of these checks from the host and others from the storage system. In some cases, you must have the Host Utilities SAN Toolkit installed before you perform the check. For example, the SAN Toolkit contains both the `basic_config` command and the `san lun` command.
- ◆ To make sure you have the current version of the system components, see the *Compatibility and Configuration Guide for NetApp FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>. Support for new components is added on an on-going basis. This online document contains a complete list of supported HBAs, platforms, applications, and drivers. These pages are updated on a regular basis to ensure that they always have information on the most current components.

Note

The examples that follow were accurate at the time this document was produced. They are not intended to show all the options available.

Remove earlier versions of the software

If you currently have the FCP Solaris Host Utilities 4.0 or FCP Solaris Attach Kit installed, you must remove it. See “[Removing an FCP Solaris Host Utilities or Attach Kit installation](#)” on page 31.

Check the version of your host operating system

Make sure you have the correct version of the operating system.

Use the `cat /etc/release` command to display information about your operating system.

Example: The following is an example of checking your operating system version on a SPARC system.

```
$ cat /etc/release
Solaris 10 11/06 s10s_u3wos_10 SPARC
Copyright 2006 Sun Microsystems, Inc. All Rights Reserved.
Use is subject to license terms.
Assembled 14 November 2006
```

Check the HBA

Supported HBAs should be installed before you install the Host Utilities. If you want to confirm that you have the supported HBAs, install the Host Utilities SAN Toolkit software package. Then use its `san lun` command to display information about HBAs on your system.

Example: The following is an example of checking your HBA using the `san lun` command.

```
# san lun fcp show adapter -v

adapter name:      lpfc1
WWPN:              10000000c94bfb30
WWNN:              20000000c94bfb30
driver name:       lpfc
model:             LPe11002-M4
model description: Emulex LPe11002-M4 4Gb 2port FC: PCIe SFF HBA
serial number:     VM54033886
hardware version:  2057706d
driver version:    6.02h; HBAAPI(I) v2.0.e, 11-07-03
firmware version:  2.50A6 (Z2F2.50A6)
Number of ports:   1
port type:         Fabric
port state:        Operational
supported speed:   4 GBit/sec
negotiated speed:  4 GBit/sec
OS device name:    /devices/pci@1e,600000/pci@0/pci@8/lpfc@0

adapter name:      lpfc2
WWPN:              10000000c94bfb31
WWNN:              20000000c94bfb31
driver name:       lpfc
model:             LPe11002-M4
model description: Emulex LPe11002-M4 4Gb 2port FC: PCIe SFF HBA
serial number:     VM54033886
hardware version:  2057706d
driver version:    6.02h; HBAAPI(I) v2.0.e, 11-07-03
firmware version:  2.50A6 (Z2F2.50A6)
Number of ports:   1
```

```

port type:          Fabric
port state:         Operational
supported speed:    4 GBit/sec
negotiated speed:   4 GBit/sec
OS device name:     /devices/pci@1e,6000000/pci@0/pci@8/lpfc@0,1

```

Check the Veritas Volume Manager

To quickly check the VxVM disks, enter the `vxdisk list` command.

Example: The following is an example of using the `vxdisk list` command. This example has been truncated to make it easier to read.

```

# vxdisk list
DEVICE      TYPE          DISK      GROUP      STATUS
Disk_0      auto:none      -         -          online invalid
Disk_1      auto:none      -         -          online invalid
FAS30200_0  auto:cdsdisk   -         -          online
FAS30200_1  auto:cdsdisk   -         -          online
FAS30200_2  auto:cdsdisk   -         -          online
FAS30200_3  auto:cdsdisk   -         -          online
...
#

```

If you want to check the volume information, use the `vxprint` command.

Example: The following is an example of using the `vxprint` command. This example has been truncated to make it easier to read.

```
# vxprint -g dg1
```

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTIL0	PUTIL0
dg	dg1	dg1	-	-	-	-	-	-
dm	c3t2d0s2	FAS30202_13	-	18800384	-	-	-	-
dm	c3t2d1s2	FAS30202_43	-	18800384	-	-	-	-
dm	c3t2d2s2	FAS30202_42	-	18800384	-	-	-	-
dm	c3t2d3s2	FAS30202_41	-	18800384	-	-	-	-
dm	c3t2d4s2	FAS30202_37	-	18800384	-	-	-	-
dm	c3t2d5s2	FAS30202_28	-	18800384	-	-	-	-
dm	c3t2d6s2	FAS30202_40	-	18800384	-	-	-	-
dm	c3t2d7s2	FAS30202_25	-	18800384	-	-	-	-
dm	c3t2d8s2	FAS30202_39	-	18800384	-	-	-	-
dm	c3t2d9s2	FAS30202_22	-	18800384	-	-	-	-
dm	c3t2d10s2	FAS30202_36	-	18800384	-	-	-	-
dm	c3t2d11s2	FAS30202_38	-	18800384	-	-	-	-
dm	c3t2d12s2	FAS30202_27	-	18800384	-	-	-	-
dm	c3t2d13s2	FAS30202_26	-	18800384	-	-	-	-
dm	c3t2d14s2	FAS30202_2	-	18800384	-	-	-	-
dm	c3t2d15s2	FAS30202_33	-	18800384	-	-	-	-
dm	c3t2d16s2	FAS30202_34	-	18800384	-	-	-	-
dm	c3t2d17s2	FAS30202_32	-	18800384	-	-	-	-
dm	c3t2d18s2	FAS30202_35	-	18800384	-	-	-	-
dm	c3t2d19s2	FAS30202_17	-	18800384	-	-	-	-
dm	c3t2d20s2	FAS30202_30	-	18800384	-	-	-	-
dm	c3t2d21s2	FAS30202_31	-	18800384	-	-	-	-
v	vol1	fsgen	ENABLED	94371840	-	ACTIVE	-	-
pl	vol1-01	vol1	ENABLED	94371840	-	ACTIVE	-	-
sd	c3t2d0s2-01	vol1-01	ENABLED	11796480	0	-	-	-
sd	c3t2d1s2-01	vol1-01	ENABLED	11796480	0	-	-	-
sd	c3t2d2s2-01	vol1-01	ENABLED	11796480	0	-	-	-
sd	c3t2d3s2-01	vol1-01	ENABLED	11796480	0	-	-	-
sd	c3t2d4s2-01	vol1-01	ENABLED	11796480	0	-	-	-
sd	c3t2d5s2-01	vol1-01	ENABLED	11796480	0	-	-	-
sd	c3t2d6s2-01	vol1-01	ENABLED	11796480	0	-	-	-
sd	c3t2d7s2-01	vol1-01	ENABLED	11796480	0	-	-	-
pl	vol1-02	vol1	ENABLED	94371840	-	ACTIVE	-	-
sd	c3t2d8s2-01	vol1-02	ENABLED	11796480	0	-	-	-
sd	c3t2d9s2-01	vol1-02	ENABLED	11796480	0	-	-	-
sd	c3t2d10s2-01	vol1-02	ENABLED	11796480	0	-	-	-
sd	c3t2d11s2-01	vol1-02	ENABLED	11796480	0	-	-	-
sd	c3t2d12s2-01	vol1-02	ENABLED	11796480	0	-	-	-
sd	c3t2d13s2-01	vol1-02	ENABLED	11796480	0	-	-	-
sd	c3t2d14s2-01	vol1-02	ENABLED	11796480	0	-	-	-
sd	c3t2d15s2-01	vol1-02	ENABLED	11796480	0	-	-	-
...								
#								

See your Veritas documentation for more information on working with the VxVM.

Check the storage system setup

You need to have your storage system correctly set up. For information on what needs to be done, see the list of storage system requirements in the section “[Pre-requisites](#)” on page 5.

Verify that the switch is installed and configured

If you have a fabric-attached configuration, check that the switch is set up and configured as outlined in the Installation and Setup instructions that shipped with your hardware. You should have completed the following tasks:

- ◆ Installed the switch in a system cabinet or rack.
- ◆ Checked that the switch firmware is supported with this Host Utility. See the *Compatibility and Configuration Guide for NetApp’s FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.
- ◆ Turned on power to the switch.
- ◆ Configured the network parameters for the switch, including its serial parameters and IP address.

Determine whether to use switch zoning

If you have a fabric-attached configuration, determine whether switch zoning is appropriate for your system setup.

Zoning requires more configuration on the switch, but it results in the following:

- ◆ Configuring the host is simplified.
- ◆ Information becomes more manageable. The output from the host tool `iostat` is easier to read because fewer paths are displayed.

To have a high-availability configuration, make sure that each LUN has at least one primary path and one secondary path through each switch. For example, if you have two switches, you would have a minimum of four paths per LUN.

It is recommended that your configuration have no more than eight paths per LUN. For more information about zoning, see the *Compatibility and Configuration Guide for NetApp’s FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Power up equipment in the correct order

The different pieces of hardware communicate with each other in a prescribed order, which means that problems occur if you turn on power to the equipment in the wrong order. Throughout the procedures that follow, this is the order you must use when powering on the equipment:

- ◆ Configured Fibre Channel switches
It can take several minutes for the switches to boot.
- ◆ Disk shelves
- ◆ Storage systems
- ◆ Host

Verify that the host and storage system can communicate

Once your setup is complete, make sure the host and the storage system can communicate.

You can verify that the host can communicate with the storage system by issuing a command from:

- ◆ Storage system's console
- ◆ A remote login that you established from the host

Using Diagnostic Utilities

Diagnostic utilities The SAN Toolkit includes several command-line utilities that provide diagnostic information you can use for troubleshooting.

When you run these utilities, they gather configuration information about different aspects of your system. They automatically collect information and place it in a directory. Then they create a compressed tar file of that directory. You can send this file to NetApp customer support.

Normally, you only use these utilities when requested to do so by Customer Support. The utilities are:

- ◆ `filer_info`. See “[Collecting storage system information](#)” on page 119 for more information.
- ◆ `solaris_info`. See “[Collecting Solaris information](#)” on page 121 for more information.
- ◆ `<switch>_info`, where `<switch>` is the brand of switch you are using (for example, `brocade_info`). See “[Collecting switch information](#)” on page 122 for more information.

General steps for executing the programs Although some of the scripts have arguments that you may want to use, the general steps for executing these scripts are the same.

Step	Action
1	Change to the <code>opt/NTAP/SANToolkit/bin</code> directory: <code>cd /opt/NTAP/SANToolkit/bin</code>
2	Enter the utility name. If the utility you are using includes arguments, you can enter them as well. Result: The configuration information that the utility gathers is placed in a compressed file.

Collecting storage system information

When to run the diagnostic script

If you are having a problem, Customer Support might ask you to run `filer_info` to collect storage system configuration information. You should only run this utility when requested to do so by Customer Support. To execute `filer_info`, complete the following steps.

Collecting storage system diagnostics

Complete the steps below to collect storage system information.

Step	Action
1	On the Solaris host, change to the <code>/opt/NTAP/SANToolkit/bin</code> directory: <code>cd /opt/NTAP/SANToolkit/bin</code>

Step	Action
2	At the host command prompt, enter the following command: <pre>filer_info [-d <i>path_and_file_name</i>] [-l <i>passwd</i>] -n <<i>storage_system_name</i>> -f</pre> -d enables you to specify a different path in which to place the resulting output. -l <i>passwd</i> allows you to specify a password to access the storage system. -n <<i>storage_system_name</i>> specifies the target of the command, where <i>storage_system_name</i> is the name of IP address of the storage system accessed. -f automatically overwrites an existing <<i>storage_system_name</i>> file in the default output directory. You might have an existing file in this directory if you previously ran the script. If you do not use the -f option, and you have an existing file, the <code>filer_info</code> script tells you to delete or rename the file. Result: The output of various status commands are placed in the output directory, where <i>storage_system_name</i> is replaced by the storage system name or IP address. The contents of this directory are placed in a tarred and compressed file named storage_system_<<i>storage_system_name</i>>.tar.Z.
3	Send the resulting compressed file to Customer Support for analysis.

Collecting Solaris information

When to run the diagnostic script

If you are having a problem, Customer Support might ask you to run `solaris_info` to collect configuration information about your host. You should only run this utility when requested to do so by Customer Support.

Collecting Solaris diagnostics

To execute `solaris_info`, complete the following steps:

Step	Action
1	On the Solaris host, change to the <code>/opt/NTAP/SANToolkit/bin</code> directory: <code>cd /opt/NTAP/SANToolkit/bin</code>
2	At the command prompt, enter the following command: <code>solaris_info [-d <directory>] [-n <file_name>] -f</code> -d <directory> specifies the location for the file that <code>solaris_info</code> creates. -n <file_name> specifies the name for the output file and directory names. -f automatically overwrites an existing .tar.Z file in the default output directory. You might have an existing file in this directory if you previously ran the script. If you do not use the -f option, and you have an existing file, the script tells you to delete or rename the file. Result: When you run <code>solaris_info</code> , the utility gathers information about your Solaris configuration and places it in a .tar.Z file. It also creates a directory and populates it with files containing output from numerous commands run on the Solaris system and contents of the <code>/var/adm/message</code> file.
3	Send the resulting file to Customer Support for analysis.

Collecting switch information

When to run the diagnostic script

If you are having a problem, Customer Support might ask you to run the switch utility script that is provided when you install the SAN Toolkit software. A utility is provided for each switch that is supported for the FCP Solaris Host Utilities. The utility name is `<switch>_info`, where *switch* is the brand of the switch (for example, if you have a McDATA switch, the file name would be `mcdata_info`). This utility collects switch configuration information. You should only run this utility when requested to do so by Customer Support.

Collecting switch diagnostics

To execute the utility, complete the following steps.

Step	Action
1	Change to the <code>/opt/NTAP/SANToolkit/bin</code> directory: <code>cd /opt/NTAP/SANToolkit/bin</code>

Step	Action
2	At the command prompt, enter the following command: <pre><switch_brand>_info [-d <i>path_and_file_name</i>] [-l <i>user:password</i>] -n <switchname> -f</pre> <switch_brand> is the type of the switch you are using; for example, brocade for a Brocade switch. The -d <i>path_and_file_name</i> option enables you to specify a different path in which to place the resulting output. The -l <i>user:password</i> option allows the choice of a user name and password on the switch. If this option is not specified, the Brocade default settings of admin:password are used for Brocade switches. The target of the command is specified by -n <i>switchname</i>, where <i>switchname</i> is the name or IP address of the switch to be interrogated. -f automatically overwrites an existing output file in the default output directory. You might have an existing file in this directory if you previously ran the script. If you do not use the -f option, and you have an existing file, the script tells you to delete or rename the file. Result: The output of various status commands on the switch are placed in a default directory where <i>switch_brand</i> is the brand of the switch, such as Brocade, and <i>switchname</i> is replaced by the switch name or IP address. The contents of this directory are placed in a file named <switch_brand>_switch_switchname.tar.Z. Example: This example uses the brocade_info tool. The switch brcd1 is connected with the user “john”, who has the password “smith”. <pre>host1> brocade_info -d /opt/NTAP/SANToolkit/bin/ -l john:smith -n brcd1</pre> The information from the switch is placed in the brocade_switch_brcd1.tar.Z file.
3	Send the resulting file to Customer Support for analysis.

Index

Symbols

/etc/system file
variables 27

A

A/A
Active/Active paths 97
A/PG-C
Active/Passive Group Concurrent paths 97
active paths
Active/Active (A/A) 97
Active/Passive Group Concurrent paths
A/PG-c 97
Array Support Library
see ASL 92
ASL
components 92
definition 92
fresh install 93
multipathing reports 97
removing 96
Attach Kits
previous kits 2

B

basic config command
description 28
options 29
basic_config
configuration script 3
boot data
copying 72
bootable LUN
creating 66
partitioning 66
bootblk
installing 72

C

checklists
installation tasks 5
commands
vxdctl enable 43
vxdiskadm 43
vxdsk list 114
vxprint 114
configuring system parameters 28
create_binding.pl
definition 46
prerequisites 48
rules for using 46
using 50

D

diagnostic utilities
<switch>_info 118
collecting Solaris information 121
collecting storage system information 119
collecting switch information 122
executing 118
filer_info 118
solaris_info 118
switch_info 122
disk shelves
powering up 117
documentation
obtaining 24
driver software
adding packages 16
downloading 13
obtaining 13

E

Emulex HBA
adding software packages 16
downloading software 13
driver software 13
firmware versions 18

- equipment
 - powering up 117

F

- FCP Solaris Attach Kit
 - removing 31
- FCP Solaris Host Utilities
 - configuration overview 34
 - contents 2
 - defined 2
 - documentation 24
 - obtaining 23
 - SAN Toolkit 3
- filer_info
 - diagnostic script 3
- firmware
 - upgrading 19
 - version 18
- format utility
 - using to label LUNs 42

H

- HBA
 - checking 113
- host
 - communicating with storage system 117
 - powering up 117

I

- igroup
 - definition 34

L

- lpfc.conf file
 - variables 26
- lputil
 - creating persistent bindings 38
- LUNs
 - discovering 40
 - labeling 42

M

- mpxio_set
 - configuration script 3
- multipathing
 - displaying information 106

O

- OpenBoot
 - definition 75
 - modifying 75

P

- persistent binding
 - creating 46
 - creating with create_binding.pl 50
 - creating with lputil 38
 - definition 34, 37
 - methods for creating 37
- powering up equipment 117

R

- removing
 - attach kit 31
 - FCP Solaris Host Utilities 31

S

- SAN boot
 - configuring 66
 - copying the boot data 72
 - definition 56
 - installing the bootblk 72
 - modifying OpenBoot 75
 - partitioning the bootable LUN 66
 - pre-requisites 57
- SAN Toolkit 3
 - components 3
 - san_version command 3
 - sanlun utility 3
- san_version
 - command 3
- sanlun

- part of SAN Toolkit 3
- syntax 107
- using 103
- scripts
 - basic_config 3
 - filer_info 3
 - mpxio_set 3
- sd.conf file
 - configuring to discover LUNs 40
 - definition 34
- single image mode
 - changing 82
 - configuration overview 85
 - definition 82
 - preparing the host 85
 - pre-requisites 84
 - reconfiguring the host 88
 - Solaris requirements 82
 - upgrading the storage system 86
 - validating zoning 87
- Solaris
 - displaying HBA information 108
 - displaying LUN information 104
 - preparing for upgrade to single image mode 85
 - requirements for single image mode 82
- storage system
 - communicating with host 117
 - powering up 117

- switch_info
 - executing 122
- switches
 - powering up 117
 - verifying setup 116
 - zoning 116
- system parameters
 - configuring 28
 - configuring manually 26
 - pre-requisites 25

U

- uninstall script 31

V

- Veritas Volume Manager
 - check VxVM 114
 - configuring LUNs 43
- Volume Manager
 - VxVM 114

Z

- zoning 116
 - validating for single image mode 87

