

FC Solaris™ Host Utilities 4.3 for Veritas™ Dynamic Multipathing

Installation and Setup Guide

NetApp, 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 215-03970_A0
June 2008

Copyright and trademark information

Copyright information

Copyright © 1994–2008 NetApp, 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 NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP 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.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. 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 NetApp.

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

Manager; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, 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, 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.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks. NetApp, Inc. NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	Host Utilities Installation and Setup Overview	1
	Host Utilities Overview.	2
	FC Solaris Host Utilities installation overview.	5
	Where to go for more information	9
Chapter 2	Managing the Driver Software	13
	Before you begin working with the drivers.	14
	Working with the Emulex LPFC software	18
	Working with the Sun Drivers for Emulex HBAs	23
	Working with the Sun Drivers for QLogic HBAs	26
Chapter 3	Configuring the Solaris Host	31
	Installing the FC Solaris Host Utilities	32
	Completing the installation and configuring system parameters.	35
	Overview of using the basic_config command.	37
	Configuring parameters on systems using LPFC drivers	38
	Configuring parameters on systems using Sun native drivers	43
	Upgrading the FC Solaris Host Utilities	46
	Removing an FC Solaris Host Utilities or Attach Kit installation	47
Chapter 4	Completing the Configuration	51
	Configuration overview.	52
	Mapping igroups and LUNs	54
	(LPFC drivers) Creating persistent bindings	55
	(LPFC drivers) Configuring the host to discover LUNs	58

	(Sun native drivers) Configuring the host to discover LUNs.	60
	Labeling the LUN and configuring Volume Management software	61
Appendix A	Installing the ASL and APM	65
	Overview of the Array Support Library	66
	Installing the ASL	68
	Example of uninstalling and installing the ASL and APM.	72
	Obtaining multipathing information	76
	Optimizing the Veritas DMP Restore Daemon.	80
	Understanding ASL error message types.	82
Appendix B	Using the create_binding.pl Script	83
	The create_binding.pl script	84
	Before you begin	86
	Using the create_binding.pl script	88
Appendix C	Using sanlun	93
	Displaying host LUN information	94
	Displaying information about the storage system paths	96
	Displaying host HBA information	98
Appendix D	Configuring SAN Boot	101
	SAN boot overview.	102
	Configuration overview.	103
	(Sun native drivers) Setting up the HBA	105
	(LPFC drivers) Setting up the HBA with Emulex drivers	111
	Creating the bootable LUN.	115
	Performing a Direct Installation to a SAN Boot LUN	121
	Copy bootblks and data to a SAN boot LUN.	122
	Modifying OpenBoot.	126

Appendix E	Changing cfmode to Single Image	.135
	Understanding cfmode	.136
	Completing prerequisites	.138
	Upgrading to single image mode	.139
 Appendix F	 Tips and Troubleshooting.	 .145
	Confirming that your system is correctly set up	.146
	Using Diagnostic Utilities	.152
	Collecting storage system information	.153
	Collecting Solaris information	.155
	Collecting switch information	.156
	 Index	 .159

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 *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do> for information on the specific Solaris versions that are 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 FC to store and retrieve data on a storage system.

Terminology

This guide uses the following terms:

- ◆ NetApp storage products (controllers, 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 FC 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 FC 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 “<code>fc test</code>” 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

A caution 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 FC Solaris in the subject line.

About this chapter

This chapter provides an introduction to the FC Solaris Host Utilities. It applies to the Host Utilities 4.3 running Veritas Dynamic Multiprocessing (DMP) with either Emulex® LightPulse® Fibre Channel (LPFC) Host Bus Adapters (HBA) or Sun native drivers (Leadville).

For details on the drivers that are supported, see the NetApp Interoperability Matrix, which is available online at <http://now.netapp.com/matrix/mtx/login.do>.

This chapter contains 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
- ◆ “[FC 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

The Host Utilities enable users to connect Solaris hosts to NetApp storage systems. Previously NetApp provided FC Solaris Attach Kits to provide this support.

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

Because the FC Solaris Host Utilities supply some tools that are common to other NetApp products, component version numbers are not tied to the Host Utilities version number. Individual tools may have separate version numbers. For example, at the time this document was produced, it supported the FC Host Utilities 4.3 and was using version 3.4 of `create_binding.pl`. See the Interoperability Matrix <http://now.netapp.com/matrix/mtx/login.do> for information about supported components.

Overview of the FC Solaris Host Utilities

The Host Utilities consist of the following tools and information:

- ◆ 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 FC Solaris Host Utilities, HBA-provided tools, or a combination FC 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*, the *Host Settings Changed by the Host Utilities* document, and a *Quick Command Reference for FC Solaris Host Utilities for Native OS and Veritas*.
- ◆ The NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do> contains the most current information on the supported stacks. NetApp adds components on an ongoing basis, so this online site is updated frequently.

Software tools provided by the Host Utilities

The FC Solaris Host Utilities bundle tools into a SAN Toolkit. These tools include scripts and commands you can use to display information about LUNs, add required settings, and troubleshoot your system.

Note

This toolkit is common across all the FC Solaris Host Utilities stacks. As a result, some of its contents apply to one stack, but not another. Having a tool that does not apply to your stack does not affect performance.

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 FC driver. It provides the HBA timeout settings for both MPxIO and Veritas DMP configurations.
- ◆ `create_binding.pl` command line tool. This tool enables you to bind the WWPNs of a single storage system's active ports to a target ID. It saves these bindings in the `/kernel/drv/lpfc.conf` file.
- ◆ `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. This utility displays information about LUNs on the storage system that are available to this host.
- ◆ `ssd_config.pl` script. This is a helper script that the `basic_config` tool calls to make setting changes for the Fibre Channel and SCSI drivers.
- ◆ `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.

About the HBAs

The Host Utilities support a number of HBAs. The supported HBAs should be installed before you install the Host Utilities. Normally, they should have the correct firmware and FCode set. To determine the firmware and FCode setting on your system, run the appropriate administration tool for your HBA.

Note

For details on the specific HBAs that are supported and the required firmware and FCode values, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do>.

If the Host Utilities setup procedure differs based on the HBA model, this guide will describe the different steps. Unless otherwise noted, the information in this guide applies to all supported versions HBA models.

Multipathing

The FC Solaris Host Utilities 4.3 described in this document use Veritas Storage Foundation Dynamic Multipathing (DMP).

FC Solaris Host Utilities installation overview

Overview of Installing the Host Utilities and setting up your system

Installing the FC 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.

Prerequisites

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
- ❖ Array Support Library
- ❖ Switches (if used)
- ❖ Remove any version of the Host Utilities or Host Attach Kit that is currently installed

See the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do> for current information about the system requirements.

2. Verify that your storage system is:

- ❖ Licensed correctly
- ❖ Using the correct `cfmode`
- ❖ Running the FC 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)

3. If you're using a fabric connection, verify that the switch is:

- ❖ Set up correctly

- ❖ Zoned (if appropriate for your system)

Note

For information on supported topologies, see the *iSCSI and Fibre Channel Configuration Guide*, which is online at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/fc_iscsi_config_guide.pdf.

- ❖ 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 `pkgadd -d NTAPSANTool.pkg` command to install the toolkit.
 - ❖ Configure storage-specific variables for your system in the correct files for your system environment. You can use the `basic_config` command to set the driver parameters.

LPFC drivers: `/etc/system` and `/kernel/drv/lpfc.conf` and `/kernel/drv/sd.conf`

Sun native drivers: `/kernel/drv/ssd.conf`.
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.
 - ❖ Map the LUN to the igroup.
 - ❖ **(LPFC drivers only)** Create persistent bindings using `lputil`. (You can also create persistent bindings using the NetApp-provided `create_binding.pl` script.)

(LPFC drivers only) In the `sd.conf` file, create an entry for every NetApp LUN.

- ❖ Configure the host to discover the LUNs.

(LPFC drivers) Either perform a reconfiguration reboot (`reboot - - -r`) on the host or use the `/usr/sbin/devfsadm` command.

(Native drivers) Use the `/usr/sbin/cfgadm -c configure cx` command, where `x` is the controller number of the HBA where the LUN is expected to be visible.

- ❖ Label the LUNs using the Solaris format utility (`/usr/sbin/format`).
- ❖ 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, verifying your setup before you install the Host Utilities, installing and configuring the Host Utilities, and creating igroups and LUNs:

- ◆ **(LPFC drivers only)** [“Working with the Emulex LPFC software”](#) on page 18
- ◆ **(Native drivers only)** [“Working with the Sun Drivers for Emulex HBAs”](#) on page 23 and [“Working with the Sun Drivers for QLogic HBAs”](#) on page 26
- ◆ [“Installing the ASL and APM”](#) on page 65
- ◆ [“Confirming that your system is correctly set up”](#) on page 146
- ◆ [“Installing the FC Solaris Host Utilities”](#) on page 32
- ◆ [“Completing the installation and configuring system parameters”](#) on page 35
- ◆ [“Mapping igroups and LUNs”](#) on page 54
- ◆ **(LPFC drivers only)** [“Using the create_binding.pl Script”](#) on page 83
- ◆ [“\(LPFC drivers\) Configuring the host to discover LUNs”](#) on page 58
- ◆ [“\(Sun native drivers\) Configuring the host to discover LUNs”](#) on page 60
- ◆ [“Labeling the new LUN”](#) on page 61
- ◆ [“Configuring volume management”](#) on page 63
- ◆ [“Displaying host LUN information”](#) on page 94
- ◆ [“Displaying information about the storage system paths”](#) on page 96
- ◆ [“Displaying host HBA information”](#) on page 98

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

Optional procedures

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

Task	Additional information
Configure SAN boot. SAN boot uses a SAN attached disk, like a LUN, as a boot device for a Solaris host	See Appendix D, “ Configuring SAN Boot ,” on page 101
Update the storage system cluster setting to single image mode. Note ————— This applies to systems using LPFC drivers. Systems using native drivers should already be set to single image mode. —————	See Appendix E, “ Changing cfmode to Single Image ,” on page 135.
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 152

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 guide was written. They provide 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, troubleshooting, operational considerations, and post-release developments	FC Solaris Host Utilities Release Notes 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 FC Solaris Host Utilities, which is available from the NOW site
The latest supported configurations	◆ NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do ◆ <i>System Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Supported Fibre Channel SAN topologies and connecting the storage system to the fabric or switch	◆ Switch documentation at http://now.netapp.com/NOW/knowledge/docs/san/ ◆ <i>iSCSI and Fibre Channel Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/fc_iscsi_config_guide.pdf
Changes to the host settings that are recommended by the Host Utilities	<i>Host Settings Affected by the Host Utilities</i> at http://now.netapp.com/NOW/knowledge/docs/san/

If you need information about...	See the following NetApp product documentation...
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 ◆ Index to Data ONTAP documentation at http://now.netapp.com/NOW/knowledge/docs/ontap
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 a storage system with environmental requirements	<i>Site Requirements</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
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	Refer to your HBA vendor documentation.

If you need information about...	See the following NetApp product 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 and DMP	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 provides information on working with the drivers. It contains information about the following

- ◆ Emulex LPFC drivers
- ◆ Sun drivers for Emulex
- ◆ Sun drivers for QLogic

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Before you begin working with the drivers”](#) on page 14
- ◆ [“Working with the Emulex LPFC software”](#) on page 18
- ◆ [“Working with the Sun Drivers for Emulex HBAs”](#) on page 23
- ◆ [“Working with the Sun Drivers for QLogic HBAs”](#) on page 26

Before you begin working with the drivers

Tasks to perform before you begin working with the drivers

Before you install new firmware or update FCodes, make sure you've taken care of the prerequisites:

- ◆ “[Verifying that system is set up correctly](#)” on page 14
- ◆ “[General information on getting the driver software](#)” on page 14
- ◆ Download the software for your HBA and extract the necessary files.
 - ❖ Emulex LPFC drivers and Sun-branded Emulex drivers: “[Downloading and extracting the Emulex software](#)” on page 15
 - ❖ Sun-branded QLogic drivers: “[Downloading and extracting the QLogic software](#)” on page 17

Verifying that system is set up correctly

Before you install the driver firmware, you must have set up:

- ◆ The storage system
 - ❖ Data ONTAP must be installed and running.
 - ❖ The storage system must be connected to the fabric or switch.
- ◆ The Solaris host
 - ❖ The host must be running the correct operating system and patches for the Host Utilities.
 - ❖ The HBA hardware must be installed.
 - ❖ The cables to the fabric or switch are correctly set up.

General information on getting the driver software

You can get the driver software from the company Web site for your HBA.

- ◆ Emulex LPFC HBAs: The Emulex software, including the Emulex utility programs and documentation, is available from the NetApp partner page on the Emulex site. Go to <http://www.emulex.com/ts/> and select NetApp. For information on downloading this software, see “[Downloading and extracting the Emulex software](#)” on page 15.
- ◆ Emulex HBAs with Sun native drivers: The Emulex software, including the Emulex utility programs and documentation, is available from the Solaris OS download section on the Emulex site. At the time this document was produced, that section was located at <http://www.emulex.com/support/solaris/index.jsp>. For information on

downloading this software, see “[Downloading and extracting the Emulex software](#)” on page 15.

- ◆ QLogic-branded HBAs: The QLogic SANsurfer CLI software and documentation are available on the QLogic support site under the Downloads section. Go to <http://www.qlogic.com> and select Downloads. You only need this software if you have to manipulate the FCode versions on QLogic-branded HBAs for SAN booting. For information on downloading this software, see “[Downloading and extracting the QLogic software](#)” on page 17.
- ◆ Sun-branded HBAs: You can also use certain Sun-branded HBAs. For more information on working with them, see the patch Readme file that Sun provides.

Downloading and extracting the Emulex software

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

Note

You may wish to download new firmware at the same time as you download the rest of the Emulex software. Your HBA may use an earlier version of the firmware than is recommended for the Host Utilities. To determine which version of the firmware you are currently running, see “[Determining the firmware version](#)” on page 24.

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

Step	Action
2	To download Emulex driver and firmware, go to the location on the Emulex Web site for the type of drivers you are using: ◆ For Emulex HBAs using LPFC drivers, go to http://www.emulex.com/ts/ and click the NetApp link. ◆ For Emulex HBAs using Sun native drivers, to the Solaris OS download section. At the time this document was produced, that section was located at http://www.emulex.com/support/solaris/index.jsp Follow the instructions on the download page to get the driver software and place it in the /tmp/emulex directory. To determine which drivers are supported, see the NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do.
3	On the Solaris host, change to the /tmp/emulex directory.
4	Use the <code>tar -xvf</code> command to extract the software from the files you downloaded. Note_____ If you are using Emulex Utilities for Sun native drivers, the .tar file you download contains two additional .tar files, each of which contains other .tar files. The file that contains the EMLXemlxu package for native drivers is <code>emlxu_kit-<version>-sparc.tar</code>. Example: The following show some of the commands you might use to extract the software, depending on which files you downloaded and which ones you need for your system. <pre>tar -xvf emlxu_kit-<version>-sparc.tar tar -xvf <solaris_version>-<driver_version>.tar tar -xvf <EmlxApps><version_number>-solaris.tar tar -xvf lpfc-<version_number>-sparc.tar</pre>

Step	Action
5	If you are using LPFC drivers, unzip and extract the HBAnyware software. Example: The following commands unzip and extract HBAnyware software: <pre>gunzip -d HBAnyware-<version_number>-sparc.tar.gz tar -xvf HBAnyware-<version_number>-sparc.tar</pre>

Downloading and extracting the QLogic software

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

Step	Action
1	On the Solaris host, create a download directory for the QLogic software. Enter: <pre>mkdir /tmp/qlogic</pre>
2	To download the SANsurfer CLI software, go to http://www.qlogic.com and click the Downloads link.
3	Under “OEM Models,” click Netapp.
4	Click the link for your card type.
5	Choose the latest multiflash or bios image available and save it to the /tmp/qlogic directory on your host.
6	Change to the /tmp/qlogic directory and uncompress files that contain the SANsurfer CLI software package. Enter: <pre>uncompress scli-<version>.SPARC-X86.Solaris.pkg.Z</pre>

Working with the Emulex LPFC software

Information in this section

This section deals with the Emulex LPFC software. It provides information on the following tasks:

- ◆ [“Installing the EMLXemlxu utilities”](#) on page 18
- ◆ [“Unbinding the HBA”](#) on page 19
- ◆ [“Adding the Emulex packages”](#) on page 20
- ◆ [“Determining the Emulex LPFC firmware version”](#) on page 20
- ◆ [“Upgrading the firmware”](#) on page 22

Emulex components

The Emulex software is provided as .tar files. You must download the correct .tar file for your system environment. The .tar files for Emulex HBAs with LPFC drivers .tar files contain the following components:

- ◆ The Host Bus Adapter driver software
- ◆ The HBA management suite, which includes HBAnyware and lputil
Sections of this document use 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.
- ◆ Emulex Fibre Channel Utilities (EMLXemlxu)

Installing the EMLXemlxu utilities

After you extract the EXLXemlxu utilities, follow the instructions in the Emulex *FCA Utilities Reference Manual* to install and use them.

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

You can use the `pkgadd` command to install the LPFC driver package and HBAnyware package. The installation process will prompt you for information. Answer the questions based on how your site is set up.

Determining the Emulex LPFC firmware version

If you are not using the Emulex firmware recommended for the Host Utilities, you must upgrade your firmware. To determine which version of firmware you should be using and which version you are actually using, complete the following steps:

Step	Action
1	Check the NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do to determine the current firmware requirements. Once you know the requirements, you can use <code>lputil</code> to determine if your firmware is current or needs updating.

Step	Action
2	Activate the lputil tool. Enter: <pre>/usr/sbin/lpfc/lputil</pre> Result: The software displays the Main Menu.
3	At the Enter choice prompt that follows the Main Menu, enter 2 to display adapter information. Result: The software displays the Adapter Information Menu.
4	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.
5	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.
6	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.

Upgrading the firmware

To upgrade your firmware version, complete the following steps:

Note

You can download the latest firmware from the NetApp 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 begin 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: - At the Enter choice prompt after the Firmware Maintenance Menu, enter 0 to return to the Main Menu. - At the Enter choice prompt after the Main Menu, enter 0 to exit lputil.

Working with the Sun Drivers for Emulex HBAs

Information in this section

This section deals with the Emulex-branded HBAs that use the Sun native drivers software. It provides information on the following tasks:

- ◆ “[Installing the EMLXemlXu utilities](#)” on page 18
- ◆ “[For more information on installing and using these utilities, see the Emulex FCA Utilities Reference Manual.](#)” on page 23
- ◆ “[Determining the firmware version](#)” on page 24
- ◆ “[Upgrading the FCode](#)” on page 24

Emulex components

The Emulex software is provided as .tar files. You must download the correct .tar file for your system environment. When you are using Emulex HBAs with native drivers, you need the .tar files containing the Emulex Fibre Channel Adapter (FCA) Utilities (EMLXemlXu).

You use these utilities to manage the firmware and FCode of the Emulex HBAs with native driver. To install and use these utilities, follow the instructions in the *Emulex FCA Utilities Reference Manual*.

Installing the EMLXemlXu utilities

After you extract the EMLXemlXu utilities, run the `emlXu_install` command to install the EMLXemlXu package:

```
# ./emlXu_install
```

Note

For more information on installing and using these utilities, see the *Emulex FCA Utilities Reference Manual*.

Determining the firmware version

If you are not using the Emulex firmware recommended for the Host Utilities, you must upgrade your firmware. To determine which version of firmware you are should be using and which version you are actually using, complete the following steps:

Note-Step	Action
1	Check the NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do to determine the current firmware requirements. Once you know the requirements, you can use <code>lputil</code> to determine if your firmware is current or needs updating.
2	Run the <code>emlxadm</code> utility. Enter: <code>/opt/EMLXemlxu/bin/emlxadm</code> Result: The software displays a list of available adapters.
3	Select the device you want to check. Result: The software displays a menu of options.
4	Exit the <code>emlxadm</code> utility by entering <code>q</code> at the <code>emlxadm></code> prompt.

Upgrading the FCode

Follow the steps below to upgrade the FCode for your firmware.

Note _____
Sun-branded HBAs have the proper firmware version pushed to the card by the native driver.

Step	Action
1	Run the <code>emlxadm</code> utility. Enter: <code>/opt/EMLXemlxu/bin/emlxadm</code> Result: The software displays a list of available adapters.
2	Select the device you want to check. Result: The software displays a menu of options.

Step	Action
3	At the emlxadm> prompt, enter: download_fw <filename> Result: The FCode is loaded onto the selected adapter.
4	Exit the emlxadm utility by entering q at the emlxadm> prompt.

Working with the Sun Drivers for QLogic HBAs

Information in this section

This section deals with the QLogic-branded HBA software. It provides information on the following tasks:

- ◆ “Installing the SANsurfer CLI package” on page 26
- ◆ “Determining the FCode on QLogic cards” on page 26
- ◆ “Determining the Emulex LPFC firmware version” on page 20
- ◆ “Upgrading the QLogic FCode” on page 27

Installing the SANsurfer CLI package

After you extract the QLogic files, you need to install the SANsurfer CLI package. To do this, complete the following steps.

Step	Action
1	Install the SANsurfer CLI package using the pkgadd command. Enter:. pkgadd -d /tmp/qlogic/scli-<version>.SPARC-X86.Solaris.pkg
2	From the directory where you extracted the QLogic software, unzip the FCode package. Enter: unzip <fcode_filename>.zip
3	For instructions on updating the FCode, please see the Upgrading “Qlogic Branded FCode” section of this doc.

Determining the FCode on QLogic cards

If you are not using the FCode recommended for the Host Utilities, you must upgrade it. To determine the FCode you should be using and which one you are actually using, complete the following steps:

Step	Action
1	Check the NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do to determine the current FCode requirements. Now, use scli utility to determine if your firmware is current or needs updating.

Step	Action
2	Run the scli utility. Enter: <code>/usr/sbin/scli</code> Result: The software displays a menu.
3	Select option 3 (HBA Information Menu). Result: The software displays the HBA Information Menu.
4	Select option 1 (Information). Result: The software displays a list of available ports.
5	Select the adapter port for which you want information. Result: The software displays information about that HBA port.
6	Write down the FCode version and press Return. Result: The software displays a list of available ports.
7	Repeat steps 5 and 6 for each adapter you want to query. When you have finished, select option 0 to return to the main menu. Result: The software displays the main menu.
8	To exit the scli utility, select option 13 (Quit).

Upgrading the QLogic FCode

To upgrade the QLogic FCode, complete the following steps.

Step	Action
1	Run the scli utility. Enter: <code>/usr/sbin/scli</code> Result: The software displays a menu.
2	Select option 8 (HBA Utilities). Result: The software displays a menu.

Step	Action
3	Select option 3 (Save Flash). Result: The software displays a list of available adapters.
4	Select the number of the adapter for which you want information. Result: The software displays a file name to use.
5	Enter the name of the file into which you want to save the flash contents. Result: The software backs up the flash contents and then waits for you to press Return.
6	Press Return. Result: The software displays a list of available adapters.
7	If you are upgrading more than one adapter, repeat steps 4 through 6 for each adapter. When you are done, select option 0 to return to the main menu.
8	Select option 0 to return to the main menu.
9	Select option 8 (HBA Utilities). Result: The software displays a menu.
10	Select option 1 (Update Flash). Result: The software displays a menu of update options.
11	Select option 1 (Select an HBA Port) Result: The software displays a list of available adapters
12	Select the appropriate adapter number. Result: The software displays a list of Update ROM options
13	Select option 1 (Update Option ROM). Result: The software requests a file name to use.

Step	Action
14	Enter the file name of the multiflash firmware bundle that you extracted from the file you downloaded from QLogic. The file name should be similar to q24mf129.bin Result: The software upgrades the FCode.
15	Press Return. Result: The software displays a menu of update options.
16	If you are upgrading more than one adapter, repeat steps 11 through 15 for each adapter.
17	When you have finished, select option 0 to return to the main menu.
18	To exit the scli utility, select option 13 (Quit).

Chapter overview

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

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Installing the FC Solaris Host Utilities”](#) on page 32
- ◆ [“Completing the installation and configuring system parameters”](#) on page 35
- ◆ [“Overview of using the basic_config command”](#) on page 37
- ◆ [“Configuring parameters on systems using LPFC drivers”](#) on page 38
- ◆ [“Configuring parameters on systems using Sun native drivers”](#) on page 43
- ◆ [“Upgrading the FC Solaris Host Utilities”](#) on page 46
- ◆ [“Removing an FC Solaris Host Utilities or Attach Kit installation”](#) on page 47

Installing the FC Solaris Host Utilities

Before you begin

Before you install the Host Utilities SAN Toolkit software package:

- ◆ You must verify that your host system meets requirements and is correctly set up. To determine the current hardware and software requirements for the Host Utilities, search the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/> for the latest hardware and software requirements.
- ◆ If you currently have the FC Solaris Host Utilities or FC Solaris Attach Kit installed, you must remove it. See “[Removing an FC Solaris Host Utilities or Attach Kit installation](#)” on page 47.

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. When you have installed the software, you can use the tools it provides to complete your setup and 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 for systems using LPFC drivers
- ◆ `sanlun` displays host and storage system information

You can also use tools from the HBA Management Suite to perform these tasks. If you wish to use HBA management suite tools, you do not need to install the FC Solaris Host Utilities. To configure the host using HBA management tools, refer to “[Completing the installation and configuring system parameters](#)” on page 35 to get the recommended settings.

Obtaining the Host Utilities software

The software is provided as compressed files on a CD-ROM. These files are also available for downloading from the NOW web site at <http://now.netapp.com/NOW/cgi-bin/software>. There is a compressed file for the stacks that support Veritas DMP and SPARC processors and a compressed file for the stack that supports Native OS using MPxIO with x86/64 (AMD) processors. You only need to install the version of the file that is appropriate for your system.

Note

To prevent you from accidentally installing the wrong file for your processor, the file will only install if it is the right file for your processor. If you have the wrong file, the installation ends with a message saying that it did not install the software.

Installing the Host Utilities

To install the FC Solaris Host Utilities software, complete the following steps.

Note

NetApp recommends that you do not install the FC Solaris Host Utilities 4.3 for Veritas on systems using versions of the Veritas Storage Foundation prior to 5.0.

Step	Action
1	Login to the host system as root.
2	Place the <code>Host Utilities</code> software in a directory on your host. You can either download this software from the NOW site or copy it from the CD-ROM For example, to copy the file to your host, you might enter: <pre>cp netapp_fcp_solaris_host_utilities_4_3_sparc.tar.gz /<directory_name></pre> Note_____ The actual file names for the Host Utilities software may be slightly different from the ones shown in the examples for these steps. _____ Example: To install the <code>netapp_fcp_solaris_host_utilities_4_3_sparc.tar.gz</code> file on a SPARC system, you might put it in the <code>/tmp</code> directory on your Solaris host. You can use following command lines to place the file in that directory and then move to the directory: <pre># cp netapp_fcp_solaris_host_utilities_4_3_sparc.tar.gz /tmp # cd /tmp</pre>

Step	Action
3	Unzip the file. Enter: <pre>gunzip netapp_fcp_solaris_host_utilities_4_3_sparc.tar.gz</pre> Result: The software unzips the tar.gz files Example: The following example uses gunzip to unzip the files for a SPARC system. <pre># gunzip netapp_fcp_solaris_host_utilities_4_3_sparc.tar.gz tar -xvf -x ./NTAPSANTool.pkg, 1199616 bytes, 2343 tape blocks</pre>
4	Untar the file. Change to the directory where you placed the file and enter: <pre>tar -xvf netapp_fcp_solaris_host_utilities_4_3_sparc.tar</pre> Result: The Host Utilities scripts are extracted to the default directory. Example: The following example uses tar -xvf to extract the Solaris installation package for a SPARC system. <pre># tar -xvf netapp_fcp_solaris_host_utilities_4_3_sparc.tar</pre>
5	Add the packages that you extracted from .tar file to your host. Enter: <pre>pkgadd -d NTAPSANTool.pkg</pre> Result: The packages are added to the opt/NTAP/SANToolkit/bin directory.

Obtaining documentation

Documentation for the Host Utilities is available on the NOW site and on the FC Solaris Host Utilities 4.3 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

Completing the Host Utilities setup

To complete the installation, you must configure the system parameters. The tasks you perform vary slightly depending on your driver.

- ◆ For LPFC drivers, see [“Configuring parameters on systems using LPFC drivers”](#) on page 38 for information on your configuration options and recommended values.

To configure these parameters, you must modify the `/etc/system` file and `/kernel/drv/lpfc.conf` and `/kernel/drv/sd.conf` files.

- ◆ For Sun native drivers, see [“Configuring parameters on systems using Sun native drivers”](#) on page 43 for information on configuring the system parameters.

To configure these parameters, you must modify the `/kernel/drv/ssd.conf` file.

There are two ways to modify these files:

- ◆ Manually editing them.
- ◆ Use the `basic_config` command to modify them. This command is provided as part of the FC Solaris Host Utilities and automatically sets these files to the correct values.

Note

The `basic_config` command does not modify the `/kernel/drv/sd.conf` file unless you are using an x86/x64 (AMD) processor with MPxIO. For information on that configuration, see the *FC Solaris Host Utilities 4.3 for Native OS with MPxIO Installation and Administration Guide*.

For more information on the `basic_config` command, see [“Overview of using the basic_config command”](#) on page 37.

Note

For a complete list of the host parameters that NetApp recommends you change and an explanation of why that 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/>).

Before you begin configuring the Host Utilities

Complete the following prerequisites before you configure system parameters.

- ◆ (LPFC drivers only) Enter the `fcv show cfmodes` command, on the storage system, to identify the storage system's cfmodes setting. Required values in the host configuration files vary, depending on the cfmodes setting.

If you plan to change the storage system cfmodes to single image, first complete the procedure in Appendix E, “[Changing cfmodes to Single Image](#),” on page 135. Then follow the procedure below to configure system parameters.

- ◆ Create your own backup of the files you are modifying:

For systems using LPFC drivers, make backups of these files:

- ❖ `/etc/system` file
- ❖ `/kernel/drv/lpfc.conf` for systems
- ❖ `/kernel/drv/sd.conf` for systems

For systems using Sun native drivers, make a backup for this file:

- ❖ `/kernel/drv/ssd.conf` for systems

The `basic_config` command automatically creates backups for you, but you can only revert to those backups once. By manually creating the backups, you can revert to them as needed.

Overview of using the `basic_config` command

Using the `basic_config` command

The `basic_config` command is a configuration script that you can use to automatically set system values to the ones recommended for use with the Host Utilities.

You execute the `basic_config` command from the command line. The options you supply with this command depend on your system setup:

- ◆ Whether your system uses Emulex's LPFC drivers or Sun's native drivers
- ◆ LPFC drivers:
 - ❖ The `cfmode`
 - ❖ The topology

Tasks performed by the `basic_config` command

You can use the `basic_config` command to set system values to the ones recommended for use with the Host Utilities.

- ◆ For systems using LPFC drivers, the `basic_config` command can make the following modifications:
 - ❖ Update items in `/kernel/drv/lpfc.conf` and `/kernel/drv/sd.conf` to the values recommended for NetApp storage systems using LPFC drivers.
 - ❖ Update timeout parameters in the `/etc/system` and `/kernel/drv/lpfc.conf` file, based on the `cfmode` that you specify.
 - ❖ Identify your Solaris OS version and set recommended values for that version in the `/etc/system` file.

See [“Configuring parameters on systems using LPFC drivers”](#) on page 38.

- ◆ For systems using Sun native drivers, the `basic_config` command can update items in `/kernel/drv/ssd.conf` to the values recommended for NetApp storage systems using Sun native drivers. See [“Configuring parameters on systems using Sun native drivers”](#) on page 43.

On systems using LPFC drivers, 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` or `/kernel/drv/ssd.conf` files..

Configuring parameters on systems using LPFC drivers

Tasks involved in configuring systems using LPFC drivers

If your system uses LPFC drivers, you must complete the configuration by performing the following tasks:

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

You can manually set the values in these files or you can set them automatically by running the basic_config command.

The sections that follow describe these tasks.

Identifying the cfmode

The values you supply for certain parameters vary depending on the storage system’s cfmode. To identify the cfmode, complete the following steps.

Step	Action
1	Access the storage system. You can do this by using a command such as telnet or going directly to the storage system.
2	On the storage system enter the following command: fcpl show cfmode Result: The storage system displays its cfmode.

Setting lpfc.conf variables manually on systems using LPFC drivers

You can manually set the /kernel/drv/lpfc.conf variables so that they match the storage system cfmode. To do this, determine the storage system’s cfmode (see “[Identifying the cfmode](#)” on page 38). Now update the file by entering the appropriate values:

- ◆ 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 manually on systems using LPFC

You can manually set the /etc/system variables so that they match the storage system cfmode. To do this, determine the storage system's cfmode (see [“Identifying the cfmode”](#) on page 38). Now update the file by entering 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 FC Solaris Host Utilities, set the following:

- ◆ set sd:sd_retry_count=5

Using basic_config to set the options automatically on systems using LPFC drivers

On systems using LPFC drivers, the basic_config command has the following format:

```
basic_config -m [cfmode] -p -r -l -s -v
```

The following table explains these options. For an example of using the basic_config command, see [“Example of using basic_config on systems with LPFC drivers”](#) on page 41.

Option	Description	Notes
-h, -help	Displays the usage page.	~
-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.
-m <cfmode>	Specifies the cfmode. Arguments are: ◆ single_image ◆ partner ◆ dual_fabric ◆ standby ◆ mixed ◆ stand_alone	To determine the cfmode, perform the following steps: 1. Connect to the storage system. You can use the telnet command to do this. 2. Execute the fcp show cfmode command.
-p	Prints out the existing settings.	~
-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.

Option	Description	Notes
-s	For switched configurations, sets the topology parameter in the /kernel/drv/lpfc.conf file to switch point-to-point (topology=2).	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.
-v	Displays version information for the basic_config script.	~

Example of using basic_config on systems with LPFC drivers

The following example steps you through the process of using basic_config to configure system parameters on a system that is using Emulex LPFC HBAs.

Step	Action
1	Check the current values for the parameters. Enter the following: <pre>basic_config -p</pre> Result: The software displays the current setting for NetApp required items in the /etc/system and /kernel/drv/lpfc.conf file.

Step	Action
2	Use the <code>basic_config -m</code> command line to configure <code>/etc/system</code> and <code>/kernel/drv/lpfc.conf</code> file items to the recommended setting. The sample command line shown below uses a <code>cfmode</code> of <code>single image</code>. basic_config -m single_image Result: The software displays recommended settings and prompts you for confirmation. You must enter y at the prompt that appears to confirm the changes and update the files. Example: This example uses the <code>basic_config -m single_image</code> command. The command displays the variables it will reset. Then it asks you to confirm this action. <pre>basic_config -m single_image Verified Emulex driver is loaded. Continuing configuration... /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? (y/yes/n/no) y Please reboot your host for the changes to take effect.</pre>

Configuring parameters on systems using Sun native drivers

Tasks involved in configuring systems using native drivers

If your system uses Sun native drivers, you must complete the configuration by setting the `/kernel/drv/ssd.conf` parameters to required values.

You can manually set the values in this file or you can set it automatically by running the `basic_config` command.

The sections that follow describe these tasks.

Setting `ssd.conf` variables manually on systems using native drivers

If your system uses Sun native drivers, you need to modify the `/kernel/drv/ssd.conf` file. The required values are:

Note

Versions of the Host Utilities using native drivers always use single image `cfmode`. You do not need to check the `cfmode`.

- ◆ `throttle_max=8`
- ◆ `not_ready_retries=300`
- ◆ `busy_retries=30`
- ◆ `reset_retries=30`
- ◆ `throttle_min=2`

You must also set the VIP/PID information to “NETAPP LUN”.

You can use the `basic_config` command to configure these parameters. It uses the following bit mask settings.

```
ssd-config-list="NETAPP LUN", "netapp-ssd-config";
netapp-ssd-config=1,0x9007,8,300,30,0,0,0,0,0,0,0,0,0,30,0,0,2,0,0;
```

Using `basic_config` to set the options on systems using native drivers and Veritas 5.0

On systems using Sun native drivers and Veritas Storage Foundation 5.0, the `basic_config` command has the following format:

```
basic_config [-vxdmp_set |
-vxdmp_unset | -vxdmp_set_force | -vxdmp_unset_force]
```

The following table explains these options. For an example of using the `basic_config` command, see [“Example of using `basic\_config` with native drivers”](#) on page 44.

Option	Description	Notes
-h, -help	Displays the usage page.	~
-p	Prints out the existing settings.	~
-v	Displays version information for the basic_config script.	~
-vxdmp_set	Adds the required settings to <code>ssd.conf</code> for VxDMP with native drivers.	For systems using DMP.
-vxdmp_set_force	Performs the same tasks as the <code>-vxdmp_set</code> option, but suppresses the confirmation prompt.	For systems using DMP.
-vxdmp_unset	Removes the entries added by the <code>-vxdmp_set</code> option.	For systems using DMP. Important: The <code>basic_config</code> script allows you to revert only one time. To revert more than one time, you must create your own backup copy of the <code>/kernel/drv/ssd.conf</code> file.
-vxdmp_unset_force	Performs the same tasks as the <code>-vxdmp_unset</code> option, but suppresses the confirmation prompt.	For systems using DMP.

Example of using basic_config with native drivers

The following example illustrates how to use `basic_config` to configure system parameters on a system that is using Sun native drivers.

Step	Action
1	Use the <code>basic_config -vxdmp_set</code> command to configure the appropriate parameters in the <code>/kernel/drv/ssd.conf</code> file to the recommended values. Example: The following example shows the type of output produced by the <code>basic_config -vxdmp_set</code> command when you execute it on a SPARC system. <pre># /opt/NTAP/SANToolkit/bin/basic_config -vxdmp_set W A R N I N G This script will modify /kernel/drv/ssd.conf to add settings required for your storage system. Do you wish to continue (y/n)?-->y</pre>

Upgrading the FC Solaris Host Utilities

Upgrading the Host Utilities or reverting to another version

You can upgrade to a newer version of the Host Utilities or revert to a 4.x or 3.0 version of the Host Utilities without any effect on system I/O.

To upgrade or revert, complete the following steps:

Step	Action
1	Use the Solaris <code>pkgrm</code> command to remove the Host Utilities software package you no longer need. For more information on removing the Host Utilities, see “Removing an FC Solaris Host Utilities or Attach Kit installation” on page 47. Note _____ Removing the software package does not remove or change the system parameter settings for that I/O stack. To remove the settings you added when you configured the Host Utilities, you must perform additional steps. You do not need to remove the settings if you are upgrading the Host Utilities. _____
2	Use the Solaris <code>pkgadd</code> command to add the appropriate Host Utilities software package. For steps on installing and configuring the Host Utilities, see “Installing the FC Solaris Host Utilities” on page 32 and “Completing the installation and configuring system parameters” on page 35.

Removing an FC Solaris Host Utilities or Attach Kit installation

Methods for removing the Host Utilities or Attach Kit

You should never install the Host Utilities on a host that currently has a version of the FC Solaris Host Utilities or Attach Kit installed. You must remove the old version before installing a new version or reverting to a different version.

The method you use for uninstalling the Host Utilities or Attach Kit depends on the version of the kit that you installed.

- ◆ For FC Solaris Host Utilities 4.x, 3.0, use the `pkgrm` command to remove the Host Utilities software package.
- ◆ For FC Solaris Attach Kit 2.0, use the `uninstall` script from the FC Solaris Attach Kit 2.0 software to uninstall the Attach Kit software.

Removing FC Solaris Host Utilities 4.x, 3.0

To remove the FC Solaris Host Utilities 4.x, 3.0, complete the following steps.

Step	Action
1	Use the following command to remove the FC Solaris Host Utilities software from the <code>/opt/NTAP/SANToolkit/bin</code> directory: <pre>pkgrm NTAPSANTool</pre> Example: The following command line removes the Host Utilities software package. <pre># pkgrm NTAPSANTool</pre>

Step	Action
2	Remove the Fibre Channel parameters that were set when you ran the <code>basic_config</code> command or that you manually set. To restore the values, you can do one of the following: ◆ Replace the system files with the backups you made before changing the values ❖ LPFC drivers: replace <code>/etc/system</code>, <code>/kernel/drv/lpfc.conf</code>, and <code>/kernel/drv/sd.conf</code>. ❖ Sun native drivers: replace <code>/kernel/drv/ssd.conf</code>. ◆ Use the <code>basic_config</code> command to revert to the saved values (you can only do this once). ❖ LPFC drivers: use the <code>basic_config -r</code> command. ❖ Sun native drivers: use the <code>basic_config -ssd_unset</code> command.

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 FC Solaris Attach Kit 2.0 software. By default, the FC Solaris Attach Kit is installed in <code>/opt/NTAP/sanlun/bin</code> .
3	From the <code>/opt/NTAP/sanlun/bin</code> directory, enter the following command to remove the existing software: <pre>./uninstall</pre> Note _____ The uninstall script automatically creates a backup copy of the <code>/kernel/drv/lpfc.conf</code> and <code>sd.conf</code> files as part of the uninstall procedure. It is a good practice, though, to create a separate backup copy before you begin the 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 52
- ◆ [“Mapping igroups and LUNs”](#) on page 54
- ◆ [“\(LPFC drivers\) Creating persistent bindings”](#) on page 55
- ◆ [“\(LPFC drivers\) Configuring the host to discover LUNs”](#) on page 58
- ◆ [“\(Sun native drivers\) Configuring the host to discover LUNs”](#) on page 60
- ◆ [“Labeling the LUN and configuring Volume Management software”](#) on page 61

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. (LPFC drivers only) 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 56 for information about creating bindings with <code>lputil</code> . Refer to “ Using the create_binding.pl Script ” on page 83 for information about creating bindings with the <code>create_binding.pl</code> script.
3. (LPFC drivers only) 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 “ (LPFC drivers) Configuring the host to discover LUNs ” on page 58 for additional information.
4. (Sun native drivers only) Display a list of controllers	Use the <code>cfgadm -al</code> command to display a list of controllers.	Refer to “ Getting the controller number ” on page 60 for additional information.

Task	Discussion	Additional information
5. Discover LUNs	(LPFC drivers) The host discovers the LUNs when you perform a reconfiguration reboot (reboot - - - r) on the host. You can also discover LUNs using the /usr/sbin/devfsadm command. (Sun native drivers) To discover the LUNs, use the command: <pre>/usr/sbin/cfgadm -c configure cx</pre> x is the controller number of the HBA where the LUN is expected to be visible	For additional information, refer to the section that applies to the HBAs your system uses: ◆ “(Sun native drivers) Configuring the host to discover LUNs” on page 60 for additional information. ◆ “(LPFC drivers) Configuring the host to discover LUNs” on page 58
6. Label LUNs	Use the Solaris format utility to label the LUNs. Note_____ Some versions of Solaris require EFI labels instead of SMI labels to avoid LUN misalignment. For the most current information on the situation and how to deal with it, see the <i>FC Solaris Host Utilities Release Notes</i>.	Refer to Solaris documentation for complete information about the format utility.
7. 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.

(LPFC drivers) 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 FC 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 56 for information about creating persistent bindings using `lputil`.
- ◆ Emulex HBAnyware — This is the 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 FC Solaris Host Utilities software.
 - ❖ To install the FC Solaris Host Utilities software, refer to “[Installing the FC Solaris Host Utilities](#)” on page 32
 - ❖ To configure persistent bindings using the `create_binding.pl` script, refer to “[Creating WWPN bindings with create_binding.pl](#)” on page 88.
- ◆ 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 FC 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 FC Solaris Host Utilities supports WWNN bindings only on storage systems 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	When you have configured the persistent bindings for all the adapters, quit lputil by performing the following steps. 1. Enter 0 to return to the Main Menu. 2. Enter 0 to exit lputil.
10	If you have not set up the sd.conf file, set it up and then reboot your system. See “Modifying the sd.conf file” on page 58.
11	To save the changes you have made, reboot the host with the reconfigure option: reboot -- -r

(LPFC drivers) 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/update_drv -f sd.conf</code> <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.

(Sun native drivers) Configuring the host to discover LUNs

Getting the controller number

Before you discover the LUNs, you need to determine what the controller number of the HBA is. To do this, complete the following step.

1	Use the <code>cfgadm -al</code> command to determine what the controller number of the HBA is. If you use the <code>/usr/sbin/cfgadm -c configure cx</code> command to discover the LUNs, you need to replace <code>x</code> with the HBA controller number. Example: To determine what the controller number of the HBA is, use the <code>cfgadm -al</code> command. <pre>\$ cfgadm -al Ap_Id Type Receptacle Occupant Condition c0 fc-fabric connected configured unknown c0::500a098187f93622 disk connected configured unknown c0::500a098197f93622 disk connected configured unknown c1 scsi-bus connected configured unknown c1::dsk/c1t0d0 disk connected configured unknown c1::dsk/c1t1d0 disk connected configured unknown c2 fc-fabric connected configured unknown c2::500a098287f93622 disk connected configured unknown c2::500a098297f93622 disk connected configured unknown</pre>
---	---

Discovering LUNs

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

Step	Action
1	To discover new LUNs, enter <code>/usr/sbin/cfgadm -c configure cx</code> where <code>x</code> is the controller number of the HBA where the LUN is expected to be visible. If you do not see the HBA in the output, check your driver installation to make sure it is correct. 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. <code>/usr/sbin/format</code> 2. At the <code>format></code> 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. The following examples show the type of output you would see on a system using LPFC drivers and on a system using Sun native drivers.

Step	Action
	Example 1: This example labels disk number 5 on a system using LPFC drivers. (Portions of this example have been removed to make it easier to review.) <pre> \$ format Searching for disks...done AVAILABLE DISK SELECTIONS: 0. c1t0d0 <SEAGATE-ST373207LSUN72G-0354 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. c7t0d0 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/sd@0,0 5. c7t0d1 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/sd@0,1 6. c7t0d2 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/sd@0,2 7. c7t0d3 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/sd@0,3 8. c7t0d4 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/sd@0,4 9. c7t0d5 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/sd@0,5 10. c7t0d6 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/sd@0,6 11. c7t0d7 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/sd@0,7 Specify disk (enter its number): 5 selecting c7t0d1 [disk formatted] ... Disk not labeled. Label it now? y \$ </pre>

Step	Action
	Example 2: This example labels disk number 2 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@8,600000/SUNW,qlc@4/fp@0,0/ssd@w500000e01008eb71,0 1. c1t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w500000e0100c6631,0 2. c6t500A098387193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 2048> /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098387193622,0 3. c6t500A098197193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 2048> /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098197193622,0 4. c6t500A098187193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 2048> /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098187193622,0 5. c6t500A098397193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 2048> /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098397193622,0 Specify disk (enter its number): 2 selecting c6t500A098387193622d0: NATE [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 `vxdctl 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

If you are using Veritas Storage Foundation for multipathing, you must install and configure the Symantec Array Support Library (ASL) and Array Policy Module (APM) for NetApp storage systems. This chapter provides information on installing and configuring the ASL and APM and how these components affect multipathing and failover.

Topics in this appendix

Topics in this appendix include the following:

- ◆ [“Overview of the Array Support Library”](#) on page 66
- ◆ [“Installing the ASL”](#) on page 68
- ◆ [“Example of uninstalling and installing the ASL and APM”](#) on page 72
- ◆ [“Obtaining multipathing information”](#) on page 76

Overview of the Array Support Library

What the Array Support Library is

The Array Support Library (ASL) is a NetApp-qualified library that provides information about storage array attributes configuration to the Device Discovery Layer (DDL) 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 Array Policy Module is

The Array Policy Module (APM) is a kernel module that defines I/O error handling, failover path selection, and other failover behavior for a specific array. The NetApp APM is customized to optimize I/O error handling and failover path selection for the NetApp environment.

What the ASL provides

The ASL provides enclosure-based naming information about SAN-attached storage systems. With enclosure-based naming, the name of the Veritas disk contains the model 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-NETAPP) disk array or an Active/Passive Concurrent (A/P-C-NETAPP) 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 NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do> for current information about the system requirements.
- ◆ Determining the ASL version if it is currently installed.
- ◆ Obtaining the ASL and the APM.
- ◆ Installing a fresh version of the ASL and APM *or* upgrading an existing version of the ASL and APM.

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 and APM software

The ASL and APM are not distributed with the Host Utilities. You obtain the ASL and APM from the Symantec web site. For Storage Foundation 5.0, the Symantec TechNote download file contains the software packages for both the ASL and the APM. You must extract these software packages and then install each one separately as described in the TechNote.

For information on the link to the ASL/APM TechNote on the Symantec Web site, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do>.

Installing the ASL and APM software you downloaded

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

Note

Before you begin, make sure you obtain the ASL/APM TechNote, which you can review at the Symantec Web site. The TechNote contains Symantec's instructions for installing the NetApp ASL and APM.

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: <code>vxddladm rmjbod vid=NETAPP</code>
3	Install the ASL and APM according to the directions provided by Symantec. The ASL/APM TechNote, which is available on the Symantec Web site, contains complete installation instructions. Note _____ You should have your LUNs set up before you install the ASL and APM. For information on setting up LUNs, see “ Completing the Configuration ” on page 51.

Step	Action
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 <code>vxddpadm listenclosure all</code> command. The output shows the model name of the storage device if you are using enclosure-based naming with VxVM. vxddpadm listenclosure all Example: The <code>vxddpadm listenclosure all</code> command shows the Enclosure Types as FAS3020 in this example. <pre># vxddpadm listenclosure all ENCLR_NAME ENCLR_TYPE ENCLR_SNO STATUS ARRAY_TYPE ===== FAS30200 FAS3020 1071852 CONNECTED A/P-C-NETAPP Disk Disk DISKS CONNECTED Disk</pre>
5	If your host is not connected to storage, use the following command: vxddladm listsupport all Example: The following is a sample of the output you see when you enter the <code>vxddladm listsupport all</code> 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 libvxsenas.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>

Step	Action																																																																																																				
6	Enter the <code>vxddmpadm listapm all</code> command to verify that the APM is installed. This command produces the following type of output. <table><tr><th>Filename</th><th>APM Name</th><th>APM Version</th><th>Array Types</th><th>State</th></tr><tr><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td></tr><tr><td>dmpaa</td><td>dmpaa</td><td>1</td><td>A/A</td><td>Active</td></tr><tr><td>dmpaaa</td><td>dmpaaa</td><td>1</td><td>A/A-A</td><td>Not-Active</td></tr><tr><td>dmpsvc</td><td>dmpsvc</td><td>1</td><td>A/A-IBMSVC</td><td>Not-Active</td></tr><tr><td>dmpap</td><td>dmpap</td><td>1</td><td>A/P</td><td>Active</td></tr><tr><td>dmpap</td><td>dmpap</td><td>1</td><td>A/P-C</td><td>Active</td></tr><tr><td>dmpapf</td><td>dmpapf</td><td>1</td><td>A/PF-VERITAS</td><td>Not-Active</td></tr><tr><td>dmpapf</td><td>dmpapf</td><td>1</td><td>A/PF-T3PLUS</td><td>Not-Active</td></tr><tr><td>dmpapg</td><td>dmpapg</td><td>1</td><td>A/PG</td><td>Not-Active</td></tr><tr><td>dmpapg</td><td>dmpapg</td><td>1</td><td>A/PG-C</td><td>Not-Active</td></tr><tr><td>dmpjbod</td><td>dmpjbod</td><td>1</td><td>Disk</td><td>Active</td></tr><tr><td>dmpjbod</td><td>dmpjbod</td><td>1</td><td>APdisk</td><td>Active</td></tr><tr><td>dmphdsalua</td><td>dmphdsalua</td><td>1</td><td>A/A-A-HDS</td><td>Not-Active</td></tr><tr><td>dmpCLARiion</td><td>dmpCLARiion</td><td>1</td><td>CLR-A/P</td><td>Not-Active</td></tr><tr><td>dmpCLARiion</td><td>dmpCLARiion</td><td>1</td><td>CLR-A/PF</td><td>Not-Active</td></tr><tr><td>dmphpalua</td><td>dmphpalua</td><td>1</td><td>A/A-A-HP</td><td>Not-Active</td></tr><tr><td>dmpnetapp</td><td>dmpnetapp</td><td>1</td><td>A/A-NETAPP</td><td>Active</td></tr><tr><td>dmpnetapp</td><td>dmpnetapp</td><td>1</td><td>A/P-C-NETAPP</td><td>Active</td></tr><tr><td>dmpnetapp</td><td>dmpnetapp</td><td>1</td><td>A/P-NETAPP</td><td>Active</td></tr></table>	Filename	APM Name	APM Version	Array Types	State	=====	=====	=====	=====	=====	dmpaa	dmpaa	1	A/A	Active	dmpaaa	dmpaaa	1	A/A-A	Not-Active	dmpsvc	dmpsvc	1	A/A-IBMSVC	Not-Active	dmpap	dmpap	1	A/P	Active	dmpap	dmpap	1	A/P-C	Active	dmpapf	dmpapf	1	A/PF-VERITAS	Not-Active	dmpapf	dmpapf	1	A/PF-T3PLUS	Not-Active	dmpapg	dmpapg	1	A/PG	Not-Active	dmpapg	dmpapg	1	A/PG-C	Not-Active	dmpjbod	dmpjbod	1	Disk	Active	dmpjbod	dmpjbod	1	APdisk	Active	dmphdsalua	dmphdsalua	1	A/A-A-HDS	Not-Active	dmpCLARiion	dmpCLARiion	1	CLR-A/P	Not-Active	dmpCLARiion	dmpCLARiion	1	CLR-A/PF	Not-Active	dmphpalua	dmphpalua	1	A/A-A-HP	Not-Active	dmpnetapp	dmpnetapp	1	A/A-NETAPP	Active	dmpnetapp	dmpnetapp	1	A/P-C-NETAPP	Active	dmpnetapp	dmpnetapp	1	A/P-NETAPP	Active
Filename	APM Name	APM Version	Array Types	State																																																																																																	
=====	=====	=====	=====	=====																																																																																																	
dmpaa	dmpaa	1	A/A	Active																																																																																																	
dmpaaa	dmpaaa	1	A/A-A	Not-Active																																																																																																	
dmpsvc	dmpsvc	1	A/A-IBMSVC	Not-Active																																																																																																	
dmpap	dmpap	1	A/P	Active																																																																																																	
dmpap	dmpap	1	A/P-C	Active																																																																																																	
dmpapf	dmpapf	1	A/PF-VERITAS	Not-Active																																																																																																	
dmpapf	dmpapf	1	A/PF-T3PLUS	Not-Active																																																																																																	
dmpapg	dmpapg	1	A/PG	Not-Active																																																																																																	
dmpapg	dmpapg	1	A/PG-C	Not-Active																																																																																																	
dmpjbod	dmpjbod	1	Disk	Active																																																																																																	
dmpjbod	dmpjbod	1	APdisk	Active																																																																																																	
dmphdsalua	dmphdsalua	1	A/A-A-HDS	Not-Active																																																																																																	
dmpCLARiion	dmpCLARiion	1	CLR-A/P	Not-Active																																																																																																	
dmpCLARiion	dmpCLARiion	1	CLR-A/PF	Not-Active																																																																																																	
dmphpalua	dmphpalua	1	A/A-A-HP	Not-Active																																																																																																	
dmpnetapp	dmpnetapp	1	A/A-NETAPP	Active																																																																																																	
dmpnetapp	dmpnetapp	1	A/P-C-NETAPP	Active																																																																																																	
dmpnetapp	dmpnetapp	1	A/P-NETAPP	Active																																																																																																	

Next steps

After you install the ASL and APM, complete the following procedures:

- ◆ If you have Data ONTAP 7.1 or higher, it is recommended to change the `cfmode` setting of your clustered systems to single image and then reconfigure your host to discover the new paths to the disk. See [“Changing cfmode to Single Image”](#) on page 135.
- ◆ On the storage system, create LUNs and map them to igroups containing the WWPNs of the host HBAs. See [“Mapping igroups and LUNs”](#) on page 54.
- ◆ On the host, discover the new LUNs and configure them to be managed by VxVM. See [“Completing the Configuration”](#) on page 51.

Example of uninstalling and installing the ASL and APM

Upgrading to a new ASL and APM when running Storage Foundation 5.0

If you are using Veritas Storage Foundation 5.0 or higher, you need to install the ASL and the APM. If you already have these installed and need to upgrade, you must first remove them.

This section contains examples that show how you remove the ASL and APM and then how you install the new version of these components. The examples assume you have already gotten the software packages from Symantec.

Uninstalling the ASL and the APM

The following is a sample uninstall of the ASL and the APM. If you were actually doing this, your output would vary slightly based on your system setup.

```
# pkginfo | grep VRTSNTAP
system      VRTSNTAPapm      Veritas NetApp Array Policy Module.
system      VRTSNTAPasl      Veritas NetApp Array Support Library
# pkgrm VRTSNTAPapm
```

The following package is currently installed:

```
VRTSNTAPapm  Veritas NetApp Array Policy Module.
              (sparc) 5.0,REV=09.12.2007.16.16
```

Do you want to remove this package? [y,n,?,q] y

```
## Removing installed package instance <VRTSNTAPapm>
```

This package contains scripts which will be executed with super-user permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] y

```
## Verifying package <VRTSNTAPapm> dependencies in global zone
```

```
## Processing package information.
```

```
## Executing preremove script.
```

```
Check if Module is loaded
```

```
## Removing pathnames in class <none>
```

```
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.9
```

```
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.8
```

```
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.10
```

```
/kernel/drv/vxapm/sparcv9 <shared pathname not removed>
```

```
/kernel/drv/vxapm/dmpnetapp.SunOS_5.9
```

```
/kernel/drv/vxapm/dmpnetapp.SunOS_5.8
```

```
/kernel/drv/vxapm/dmpnetapp.SunOS_5.10
```

```
/kernel/drv/vxapm <shared pathname not removed>
```

```

/kernel/drv <shared pathname not removed>
/kernel <shared pathname not removed>
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.9
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.8
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.10
/etc/vx/apmkey.d/64 <shared pathname not removed>
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.9
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.8
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.10
/etc/vx/apmkey.d/32 <shared pathname not removed>
/etc/vx/apmkey.d <shared pathname not removed>
/etc/vx <shared pathname not removed>
/etc <shared pathname not removed>
## Updating system information.

Removal of <VRTSNTAPapm> was successful.
#

# pkgrm VRTSNTAPasl

The following package is currently installed:
    VRTSNTAPasl   Veritas NetApp Array Support Library
                  (sparc) 5.0,REV=11.19.2007.14.03

Do you want to remove this package? [y,n,?,q] y

## Removing installed package instance <VRTSNTAPasl>

This package contains scripts which will be executed with super-user
permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] y
## Verifying package <VRTSNTAPasl> dependencies in global zone
## Processing package information.
## Executing preremove script.
Unloading the library
## Removing pathnames in class <none>
/etc/vx/lib/discovery.d/libvxnetapp.so.2
/etc/vx/lib/discovery.d <shared pathname not removed>
/etc/vx/lib <shared pathname not removed>
/etc/vx/aslkey.d/libvxnetapp.key.2
/etc/vx/aslkey.d <shared pathname not removed>
## Executing postremove script.
## Updating system information.

Removal of <VRTSNTAPasl> was successful.
#

```

Installing the ASL and the APM

The following is a sample installation of the ASL and the APM when you have Veritas Storage Foundation 5.0. If you were actually doing this, your output would vary slightly based on your system setup.

```
# pkgadd -d . VRTSNTAPas1

Processing package instance <VRTSNTAPas1> from </tmp>

Veritas NetApp Array Support Library(sparc) 5.0,REV=11.19.2007.14.03
Copyright © 1990-2006 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo are trademarks or registered trademarks of
Symantec Corporation or its affiliates in the U.S. and other countries. Other
names may be trademarks of their respective owners.

The Licensed Software and Documentation are deemed to be "commercial computer
software" and "commercial computer software documentation" as defined in FAR
Sections 12.212 and DFARS Section 227.7202.
Using </etc/vx> as the package base directory.
## Processing package information.
## Processing system information.
    3 package pathnames are already properly installed.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <VRTSNTAPas1> [y,n,?] y

Installing Veritas NetApp Array Support Library as <VRTSNTAPas1>

## Installing part 1 of 1.
/etc/vx/aslkey.d/libvxnetapp.key.2
/etc/vx/lib/discovery.d/libvxnetapp.so.2
[ verifying class <none> ]
## Executing postinstall script.
Adding the entry in supported arrays
Loading The Library

Installation of <VRTSNTAPas1> was successful.
#

# pkgadd -d . VRTSNTAPapm

Processing package instance <VRTSNTAPapm> from </tmp>
```

```
Veritas NetApp Array Policy Module.(sparc) 5.0,REV=09.12.2007.16.16
Copyright 1996-2005 VERITAS Software Corporation. All rights reserved.
VERITAS, VERITAS SOFTWARE, the VERITAS logo and all other VERITAS
product names and slogans are trademarks or registered
trademarks of VERITAS Software Corporation in the USA and/or other
countries. Other product names and/or slogans mentioned herein may
be trademarks or registered trademarks of their respective companies.
```

```
Using </> as the package base directory.
```

```
## Processing package information.
```

```
## Processing system information.
```

```
9 package pathnames are already properly installed.
```

```
## Verifying disk space requirements.
```

```
## Checking for conflicts with packages already installed.
```

```
## Checking for setuid/setgid programs.
```

```
This package contains scripts which will be executed with super-user
permission during the process of installing this package.
```

```
Do you want to continue with the installation of <VRTSNTAPapm> [y,n,?] y
```

```
Installing Veritas NetApp Array Policy Module. as <VRTSNTAPapm>
```

```
## Installing part 1 of 1.
```

```
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.10
```

```
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.8
```

```
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.9
```

```
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.10
```

```
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.8
```

```
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.9
```

```
/kernel/drv/vxapm/dmpnetapp.SunOS_5.10
```

```
/kernel/drv/vxapm/dmpnetapp.SunOS_5.8
```

```
/kernel/drv/vxapm/dmpnetapp.SunOS_5.9
```

```
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.10
```

```
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.8
```

```
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.9
```

```
[ verifying class <none> ]
```

```
## Executing postinstall script.
```

```
Installation of <VRTSNTAPapm> was successful.
```

```
#
```

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 NetApp (A/A-NETAPP)—All paths to storage are active and simultaneous I/O is supported on all paths. If a path fails, I/O is distributed across the remaining paths.
- ◆ Active/Passive Concurrent-NetApp (A/P-C-NETAPP) - The array 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 document:

- ◆ *Veritas Volume Manager Administrator's Guide* - for details about system management

How the storage system's FC 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-NETAPP; otherwise, it will be A/P-C-NETAPP.

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-NETAPP) array types.

How sanlun displays the array type

The Host Utilities SAN Toolkit includes the `san lun` utility in the `opt/NTAP/SANToolkit/bin` directory. You use the `san lun` 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/P-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>vxddmpadm getsubpaths ctlr=<controller_name></pre> controller_name is the controller displayed under “CTRL - NAME” in the output of the <code>vxddmpadm getsubpaths dmpnodename</code> 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># vxddmpadm 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.

For the most up-to-date recommendations on the values, see the General Operational Considerations page of the *FC Solaris Host Utilities Release Notes*. The Release Notes are online in the SAN/iSAN Information Library at <http://now.netapp.com/NOW/knowledge/docs/san/>.

Setting the restore daemon interval

At the time this document was prepared, NetApp recommended you set the restore daemon interval value to 60 seconds to improve the recovery of previously failed paths.

To set the `vxrestore` daemon, complete the following steps.

Step	Action
1	Stop the restore daemon. <code>/usr/sbin/vxdmpadm stop restore</code>
2	Change the restore daemon setting to 60. <code>/usr/sbin/vxdmpadm start restore interval=60 policy=check_disabled</code> Note _____ This step reconfigures and restarts the restore daemon without the need for an immediate reboot. _____

Step	Action
3	Edit the edit the /lib/svc/method/vxvm-sysboot file in Solaris 10 or the /etc/init.d/vxvm-sysboot file in Solaris 9 to make the new restore daemon interval persistent across reboots. By default, the restore daemon options are: <pre>restore_daemon_opts="interval=300 policy=check_disabled"</pre> Edit the restore daemon options so that the interval is 60 seconds: <pre>restore_daemon_opts="interval=60 policy=check_disabled"</pre>
4	Save and exit the /etc/init.d/vxvm-sysboot file.
5	Verify the changes. <pre>/usr/sbin/vxdmpadm stat restored</pre>

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.	

Appendix overview

This appendix describes how to use the `create_binding.pl` script to create persistent bindings. This script is bundled with the FC 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 56
- ◆ 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 84
- ◆ “[Before you begin](#)” on page 86
- ◆ “[Using the create_binding.pl script](#)” on page 88

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 bind the WWPNs of a single storage system's active ports to a target ID. The tool then 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 FC 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 FC 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 FC Solaris Host Utilities is installed.

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

Refer to “[Installing the FC Solaris Host Utilities](#)” on page 32 for information about downloading and installing the FC 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 `fcpl 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 FC 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 when Veritas DMP has not yet been installed. Use this option if you have multiple paths to the storage system, but have not 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 -d</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. _____

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 94
- ◆ [“Displaying information about the storage system paths”](#) on page 96
- ◆ [“Displaying host HBA information”](#) on page 98

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: ◆ host device filename is the special device file name for the disk on Solaris (which might represent a storage system LUN). ◆ all lists all storage system LUNs under /dev/rdsk. ◆ storage system name lists all storage system LUNs under /dev/rdsk on a specified storage system. ◆ storage system name:storage system pathname lists all storage system LUNs under /dev/rdsk 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 <code>sanlun show</code> command in verbose mode. <pre># ./sanlun lun show -v filer: lun-pathname device filename adapter lun size lun state filerX: /vol/vol1/hostA_lun2 /dev/rdsk/c0t500A098487093F9Dd5s2 emlxs0 3g (3221225472) GOOD Serial number: HnTMWZEqHyD5 Filer FCP nodename:500a098087093f9d Filer FCP portname:00a098397093f9d Filer adapter name: 0a Filer IP address: 10.60.181.66 Filer volume name:vol1 FSID:0x331ee81 Filer qtree name:/vol/vol1 ID:0x0 Filer snapshot name: ID:0x0 LUN partition table permits multiprotocol access: no why: bad starting cylinder or bad size for 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: <code>cd /opt/NTAP/SANToolkit/bin</code>																																			
3	At the host command line, enter the following command to display LUN information: <pre>./sanlun lun show -p all -v</pre> -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/c0t500A098397093F9Dd0s2 (c0t500A098397093F9Dd0s2) 3g (3221225472) 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/c0t500A098497093F9Dd0s2</td><td></td><td></td><td>emlxs0 0b</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c1t500A098297093F9Dd0s2</td><td></td><td></td><td>emlxs1 0d</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c0t500A098487093F9Dd0s2</td><td></td><td></td><td>emlxs0 0b</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c1t500A098287093F9Dd0s2</td><td></td><td></td><td>emlxs1 0d</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/c0t500A098497093F9Dd0s2			emlxs0 0b	A/PG-C	enabled	primary	/dev/rdsk/c1t500A098297093F9Dd0s2			emlxs1 0d	A/PG-C	enabled	secondary	/dev/rdsk/c0t500A098487093F9Dd0s2			emlxs0 0b	A/PG-C	enabled	secondary	/dev/rdsk/c1t500A098287093F9Dd0s2			emlxs1 0d	A/PG-C
path state	path type	device filename	host HBA	target port	partner port	Multipath policy																														
enabled	primary	/dev/rdsk/c0t500A098497093F9Dd0s2			emlxs0 0b	A/PG-C																														
enabled	primary	/dev/rdsk/c1t500A098297093F9Dd0s2			emlxs1 0d	A/PG-C																														
enabled	secondary	/dev/rdsk/c0t500A098487093F9Dd0s2			emlxs0 0b	A/PG-C																														
enabled	secondary	/dev/rdsk/c1t500A098287093F9Dd0s2			emlxs1 0d	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/P-C (Active/Passive 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 FC adapters. adapter name lists information for a specified adapter. Result: The FC adapter information is displayed. Example: The following command line displays information about the adapter. <pre>#./sanlun fcp show adapter -v adapter name: emlxs0 WWPN: 10000000c94be338 WWNN: 20000000c94be338 driver name: emlxs model: LP11002-M4 model description: Emulex LP11002-M4 4Gb 2port FC: PCI-X2 SFF HBA serial number: VM53824441 hardware version: 1036406d driver version: 2.20k (2007.06.04.09.35) firmware version: 2.10a10 Number of ports: 1 of 2 port type: Fabric port state: Operational supported speed: 1 Gbit/sec, 2 Gbit/sec, 4 Gbit/sec negotiated speed: 4 Gbit/sec OS device name: /dev/cfg/c0 adapter name: emlxs1 WWPN: 10000000c94be339 WWNN: 20000000c94be339 driver name: emlxs model: LP11002-M4 model description: Emulex LP11002-M4 4Gb 2port FC: PCI-X2 SFF HBA serial number: VM53824441 hardware version: 1036406d driver version: 2.20k (2007.06.04.09.35) firmware version: 2.10a10 Number of ports: 2 of 2 port type: Fabric port state: Operational supported speed: 1 Gbit/sec, 2 Gbit/sec, 4 Gbit/sec negotiated speed: 4 Gbit/sec OS device name: /dev/cfg/c1</pre>

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 NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do>.

Topics in this appendix

Topics in this appendix include the following:

- ◆ “[SAN boot overview](#)” on page 102
- ◆ “[Configuration overview](#)” on page 103
- ◆ “[\(Sun native drivers\) Setting up the HBA](#)” on page 105
- ◆ “[Creating the bootable LUN](#)” on page 115
- ◆ “[Performing a Direct Installation to a SAN Boot LUN](#)” on page 121
- ◆ “[Copy bootblks and data to a SAN boot LUN](#)” on page 122
- ◆ “[Modifying OpenBoot](#)” on page 126

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 boot LUN for both storage and boot needs, thereby consolidating and centralizing storage.

Contents of this appendix

This appendix provides steps for configuring a SAN boot LUN and installing a new OS image on it. You can apply these steps to most configurations.

Keep in mind that you can use other methods to create a SAN boot LUN. This appendix does not describe other examples of bootable LUNs, such as configurations that boot multiple volumes or use diskless servers.

If you are using Sun native drivers, refer to Solaris documentation for details about additional configuration methods. In particular, 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 procedures and examples in this chapter assume the following:

- ◆ The FC 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.
- ◆ Bootcode/Fcode is downloaded and installed on the HBA.
 - ❖ For Emulex HBAs, the FCode is available on the Emulex site. For details on how to load the FCode onto the HBA, follow the steps in [“Upgrading the FCode”](#) on page 24.
 - ❖ For QLogic-branded HBAs, the FCode is available on the QLogic site. For details on how to load the FCode onto the HBA, follow the steps in [“Upgrading the QLogic FCode”](#) on page 27.
 - ❖ For Sun-branded QLogic HBAs, the FCode is available as a patch from Sun. Follow steps in the patch Readme for details on how to load this FCode onto the HBA.
- ◆ If you are using Emulex HBAs, you must have the Emulex FCA utilities with EMLXemlxu and the EMLXdrv have been installed. For more information on these tools, see [“Installing the EMLXemlxu utilities”](#) on page 18.
- ◆ If you are using QLogic-branded HBAs, you must have the SANsurfer SCLI utility installed. For more information on this tool, see [“Installing the SANsurfer CLI package”](#) on page 26.
- ◆ If you are using the Emulex LPFC driver, make sure you have configured persistent binding. There are three basic ways to configure the bindings:
 - ❖ Manually
 - ❖ By using create_binding.pl
 - ❖ By using lputil.

For information about host, storage system and firmware versions, refer to the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>.

Configuration overview

Configuring a bootable LUN involves:

1. Making sure the HBA is set to the appropriate Bootcode

For more information on setting this, see the chapter “[Managing the Driver Software](#)” on page 13.

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. Selecting the method for installing to the SAN booted LUN
 - ❖ Directly Install the bootblks onto the bootable LUN
 - ❖ Perform a file system dump and restore to a SAN booted LUN
 1. Install the bootblks onto the bootable LUN
 2. Copy the boot data from the source disk onto the bootable LUN
4. Modifying Bootcode
 - ❖ Verify the OpenBoot version
 - ❖ Set the FC topology to the bootable LUN
 - ❖ Bind the adapter target and the bootable LUN
 - ❖ Create an alias for the bootable LUN
5. Rebooting the system.

(Sun native drivers) Setting up the HBA

Setting up the HBA

Part of configuring a bootable LUN involves setting up your HBA. You may need to shut down the system and switch the HBA mode.

Note

The steps in this section apply only to the SUN native HBAs. If you are using LPFC HBAs, see the section “(LPFC drivers) Setting up the HBA with Emulex drivers” on page 111.

Changing the HBA mode

The steps you need to perform depend on the type of HBA you are using. For information on changing the mode on:

- ◆ Emulex HBAs, see “Changing the Emulex HBA to SFS mode” on page 105
- ◆ QLogic HBAs, see “Changing the QLogic HBA to enable FCode compatibility” on page 109

Changing the Emulex HBA to SFS mode

To change the mode on an Emulex HBA from an SD compatible mode to an SFS mode, you must bring the system down and then change each HBA. Complete the following steps.

WARNING

These steps will change the device definition from lpfc@ to emlxs@. Doing this will cause the controller instance to be incremented. Any devices current on the controllers that are being modified will receive new controller numbers. If you are currently mounting these devices via the /etc/vfstab file, these entries will be invalid.

Step	Action
1	At the operating system prompt, issue the <code>init 0</code> command. <code># init 0</code>
2	When the <code>ok</code> prompt appears, enter the <code>setenv auto-boot? false</code> command. <code>ok setenv auto-boot? false</code>
3	Enter the <code>reset-all</code> command. <code>ok reset-all</code>

Step	Action
4	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 SFS mode. If the mode had been changed, the devices would have displayed <code>.../lpfc</code> instead of <code>.../emlxs</code>. See Step 5 for information setting the devices to SFS 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/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>

Step	Action
5	Select the first Emulex device and set it to SFS mode using the <code>set -sfs-boot</code> command. Doing this changes the devices to emlxs devices. Example: The first command, <code>show-devs</code>, shows the device name. The next command, <code>select</code>, selects the device <code>lpfc@0</code>. The command, <code>set -sfs-boot</code>, sets the SFS mode. Some of the output has been truncated to make this easier to read. <pre>(0) show-devs ...output truncated... /pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0,1 /pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0 ...output truncated... {0} ok select /pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0 {0} ok set-sfs-boot Flash data structure updated. Signature 4e45504f Valid_flag 4a Host_did 0 Enable_flag 5 SFS_Support 1 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. *** {0} ok</pre>
6	Repeat the previous step for each Emulex device.
7	Enter the <code>reset-all</code> command to update the devices. Example: The <code>reset-all</code> command updates the Emulex devices with the new mode. <pre>ok reset-all Resetting ...</pre>

Step	Action
8	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 the Emulex devices are showing up as emlx devices. To continue the example in Step 5, the device selected there, <code>/pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0</code>, has been changed to an emlx device. In a production environment, you would want to change all the devices to emlx. <pre> ok show-devs /pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0,1 /pci@0/pci@0/pci@8/pci@0/pci@1/emlx@0 /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/scsi@6,1 /pci@8,700000/scsi@6 /pci@8,700000/usb@5,3 </pre>
9	Set the auto-boot? back to true and boot the system with a reconfiguration boot. Example: The boot command brings the system back up. <pre> ok setenv auto-boot? true ok boot -r </pre>

Changing the QLogic HBA to enable FCode compatibility

To enable FCode compatibility on a QLogic HBA, 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. <code># init 0</code>
2	When the <code>ok</code> prompt appears, enter the <code>setenv auto-boot? false</code> command. <code>ok setenv auto-boot? false</code>
3	Enter the <code>reset-all</code> command. <code>ok reset-all</code>
4	Issue the <code>show-devs</code> command to see the current device names. Example: Execute the <code>show-devs</code> shows to see whether there is FCode compatibility. <code>ok show-devs</code> ... <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1</code> <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0</code> <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1/fp@0,0</code> <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1/fp@0,0/disk</code> <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0/fp@0,0</code> <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0/fp@0,0/disk</code>
5	Select the first QLogic device. Example: Use the <code>select</code> command to select the first QLogic device. <code>ok select /pci@7c0/pci@0/pci@8/QLGC,qlc@0,1</code> QLogic QLE2462 Host Adapter Driver(SPARC): 1.16 03/10/06
6	Execute the <code>set-mode</code> command and, if needed, set the compatibility mode to <code>fcode</code> . Example: The <code>set -mode</code> command sets the mode to <code>fcode</code> . <code>ok set-mode</code> Current Compatibility Mode: <code>fcode</code> Do you want to change it? (y/n) Choose Compatibility Mode: <code>0 - fcode</code> <code>1 - bios</code> <code>enter: 0</code> Current Compatibility Mode: <code>fcode</code>

Step	Action
7	Execute the <code>set-fc-mode</code> command and, if needed, set the Fcode Mode to <code>qlc</code>. Example: The <code>set -mode</code> command sets the mode to <code>fcode</code>. <pre>ok set-fc-mode Current Fcode Mode: qlc Do you want to change it? (y/n) Choose Fcode Mode: 0 - qlc 1 - qla enter: 0 Current Fcode Mode: glc</pre>
8	Repeat the previous steps to configure each QLogic device.
9	Enter the <code>reset-all</code> command to update the devices. Example: The <code>reset-all</code> command updates the QLogic devices with the new mode. <pre>ok reset-all Resetting ...</pre>
10	Set the <code>auto-boot?</code> back to <code>true</code> and boot the system with a reconfiguration boot. Example: The <code>boot</code> command brings the system back up. <pre>ok setenv auto-boot? true ok boot -r</pre>

(LPFC drivers) Setting up the HBA with Emulex drivers

Setting up the LPFC HBA

Part of configuring a bootable LUN involves setting up your LPFC drivers. Before you do this, you need to shut down the system and set the Emulex HBA to SD mode.

Note — The steps in this section apply only to the LPFC HBAs. If you are using Sun native HBAs, see the section “Setting up the HBA (Sun native HBAs).”

WARNING — These steps may change the device definition from `lpfc emlxs@` to `lpfc@`. Doing this will cause the controller instance to be incremented. Any devices current on the controllers that are being modified will receive new controller numbers. If you are currently mounting these devices via the `/etc/vfstab` file, these entries will be invalid.

Changing the Emulex 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. <code># init 0</code>
2	When the <code>ok</code> prompt appears, enter the <code>setenv auto-boot? false</code> command. <code>ok setenv auto-boot? false</code>
3	Enter the <code>reset-all</code> command. <code>ok reset-all</code>

Step	Action
4	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>.../emlxs lpfc</code> instead of <code>.../lpfc</code>. See Step 5 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/emlxs@3 /pci@8,700000/emlxs@1 /pci@8,700000/scsi@6,1 /pci@8,700000/scsi@6 /pci@8,700000/usb@5,3 </pre>

Step	Action
5	Select the first Emulex device and set it to SD mode using the <code>set -sd-boot</code> command. This changes the Emulex devices to lpfc devices. Example: The first command, <code>select</code>, selects the Emulex device 700000/emlxs@1. The second command, <code>set -sd-boot</code>, sets the SD mode. <pre> ok select /pci@8,700000/lpfc@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>
6	Repeat the previous step for each Emulex device.
7	Enter the <code>reset -all</code> command to update the devices. Example: The <code>reset -all</code> command updates the Emulex devices with the new mode. <pre> ok reset-all Resetting ... </pre>

Step	Action
8	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>
9	Set the auto-boot? back to true and boot the system with a reconfiguration boot. Example: The boot command brings the system back up. <pre> ok setenv auto-boot? true ok boot -r </pre>

Creating the bootable LUN

Create the LUN

Use standard storage system commands and procedures to create and map a LUN assigned to a host.

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: - Execute 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 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. c2t500A098786F7CED2d11 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b 5. c2t500A098786F7CED2d12 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,c 6. c2t500A098786F7CED2d13 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,d 7. c2t500A098786F7CED2d14 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,e 8. c2t500A098786F7CED2d15 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,f</pre> Note_____ If you want to set up additional SAN boot configurations, you can use the san lun command to find the primary path. The examples in this section use the c1t1d0 device. _____ - 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 /etc/vfstab file. Result: The system selects the Solaris boot disk and displays partition options

Step	Action
2	Look in the /etc/vfstab file and write down the information about the boot file systems. Example: The following example shows the contents of the /etc/ vfstab file. Devices c1t1d0s0 and c1t1d0s1 (in bold) refer to the source boot device and must be replicated on the bootable LUN. <pre>#vi /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/c1t1d0s1 - - swap - no - /dev/dsk/c1t1d0s0 /dev/rdsk/c1t1d0s0 / ufs 1 no - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes -</pre>

Step	Action																																																						
3	Display partition information for the current Solaris boot disk: - At the <code>format></code> prompt, enter p to display partition menu options - At the <code>partition></code> prompt, enter p to print the partition table for the current Solaris boot device. Write down the following information: ❖ 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). The devices that you must include are referred to in the <code>/etc/vfstab</code> file. 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: 6142 + 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>683 - 5461</td><td>14.00GB</td><td>(4779/0/0) 29362176</td></tr><tr><td>1</td><td>swap</td><td>wu</td><td>0 - 682</td><td>2.00GB</td><td>(683/0/0) 4196352</td></tr><tr><td>2</td><td>backup</td><td>wu</td><td>0 - 6141</td><td>17.99GB</td><td>(6142/0/0) 37736448</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> - At the <code>partition></code> prompt, enter q to go back to the top level <code>format></code> prompt. 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. _____	Part	Tag	Flag	Cylinders	Size	Blocks	0	root	wm	683 - 5461	14.00GB	(4779/0/0) 29362176	1	swap	wu	0 - 682	2.00GB	(683/0/0) 4196352	2	backup	wu	0 - 6141	17.99GB	(6142/0/0) 37736448	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	683 - 5461	14.00GB	(4779/0/0) 29362176																																																		
1	swap	wu	0 - 682	2.00GB	(683/0/0) 4196352																																																		
2	backup	wu	0 - 6141	17.99GB	(6142/0/0) 37736448																																																		
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																																																		

Step	Action
4	Use format to create the same size root and swap slice for the SAN boot disk. 1. At the format> prompt, type disk to get a list of disks. 2. At the specify disk prompt, enter the number that corresponds to the LUN mapped to host controller and target0 Note_____ Make a note of the identifying information for the LUN you select. You will use this information in “Modifying OpenBoot” on page 126, when you bind the bootable LUN to the adapter and create a boot alias. Pay close attention to the “cXtYdZ” information. Example: This example selects c2t500A098786F7CED2d11 as the bootable LUN. <pre>format> disk 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. c2t500A098786F7CED2d11 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b 5. c2t500A098786F7CED2d12 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,c 6. c2t500A098786F7CED2d13 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,d 7. c2t500A098786F7CED2d14 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,e 8. c2t500A098786F7CED2d15 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,f Specify disk (enter its number): Specify disk (enter its number): 4 selecting c2t500A098786F7CED2d11 [disk formatted]</pre>

Step	Action
5	Display partition information for the current 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.
6	At the partition> prompt, enter the number of the slice you want to modify to match the host boot device. Repeat this step for each slice. Set any unused slice to 0 cylinders in length.
7	At the partition> prompt, enter label l to write the new label to the disk.
8	Exit the format utility. 1. At the partition> prompt, enter q. 2. At the format> prompt, enter q.

Performing a Direct Installation to a SAN Boot LUN

Installing directly onto a SAN boot LUN

If you are using Sun native drivers, you can perform a direct installation onto the SAN boot LUN. You can install the operating system using:

- ◆ A CD/DVD installation
- ◆ A jumpstart installation

During the installation, select the device you created in the steps in the section [“Creating the bootable LUN”](#) on page 115.

Copy bootblks and data to a SAN boot LUN

- 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/`uname -i`/lib/fs/ufs/bootblk /dev/rdisk/<device_name></pre> Example: The following example installs bootblks in the <code>/usr/platform/'uname -i'</code> directory on <code>/dev/rdisk/c2t500A098786F7CED2d11s0</code>. <pre>#/usr/sbin/installboot /usr/platform/sun4u/lib/fs/ufs/bootblk /dev/rdisk/c2t500A098786F7CED2d11s0</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 file system.
 - ◆ 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: newfs <device> Example: The following creates a USF file system at /dev/rdisk/c2t500A098786F7CED2d11s0. <pre>#newfs on /dev/rdisk/c2t500A098786F7CED2d11s0 newfs: /dev/rdisk/c2t500A098786F7CED2d11s0 last mounted as / newfs: construct a new file system /dev/rdisk/c2t500A098786F7CED2d11s0: (y/n)? y Warning: 4096 sector(s) in last cylinder unallocated /dev/rdisk/c2t500A098786F7CED2d11s0: 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/c2t500A098786F7CED2d11s1 - - swap - no - /dev/dsk/c2t500A098786F7CED2d11s0 /dev/rdisk/c2t500A098786F7CED2d11s0 / 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: mount <filesystem_name> /<mountpoint> Example: The following mounts the file system on the bootable LUN under /mnt. <pre>#mount /dev/dsk/c2t500A098786F7CED2d11s0 /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 restores it to the SAN boot disk. <pre>#init S #df -k Filesystem kbytes used avail capacity Mounted on /dev/dsk/c1t1d0s0 10327132 3882023 6341838 38% / # ufsdump 0f - /dev/dsk/c1t1d0s0 (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/dsk/c1t1d0s0 (sunv440-shu05:/) 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>

Step	Action
4	Edit the /etc/vfstab file on the SAN boot LUN to refer to the correct device. You should use a primary to the LUN. To determine the primary path, run <code>san lun show -p</code> refer to the bootable LUN. The <code>san lun</code> command comes with the Host Utilities. 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/c2t500A098786F7CED2d11s1 - - swap - yes - /dev/dsk/c2t500A098786F7CED2d11s0 /dev/rdsk/c2t500A098786F7CED2d11s0 / ufs 1 yes - /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 host uses to start up the system. OpenBoot firmware also includes the hardware level user interface that you use to configure the bootable LUN.

Modifying OpenBoot

Modifying OpenBoot involves creating an alias for the bootable LUN. This alias substitutes for the device address (/pci@8,700000/emlxs@1/sd@3,0 in the example that follows) during subsequent boot operations.

To modify OpenBoot, complete the following steps.

Step	Action
1	Activate the OpenBoot environment by performing the following steps: 1. Halt the Solaris operating system by entering the command: halt Result: The Open Boot program displays the OpenBoot prompt. 2. Stop automatic reboot during OpenBoot configuration by entering the command: setenv auto-boot? false 3. Reset the host and activate your changes by entering: reset-all
2	Check the OpenBoot version that is currently installed to be sure the host is running the most recent OpenBoot firmware (OBP). Enter the command: .version

Step	Action
3	Run the show-disks command to get the correct device path name for your Fibre Channel card. It is important to look for the device path associated with the LUN you used during the earlier format step. Due to the way the OBP sees hardware, the device may show up with either “disk” or “ssd” as the terminating characters and those characters may differ from what was seen in the format output. Example: This output shown here continues the previous example and uses the information that was written down during the format command. The LUN is “w500a098786f7ced2,b” and the device path is “/pci@1d,700000/emlx@2/fp@0,0/ssd.” <pre>4. c2t500A098786F7CED2d11 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b</pre> The following example simply shows the type of output you see when you execute the show-disks command. <pre>{1} ok show-disks a) /pci@1f,700000/scsi@2,1/disk b) /pci@1f,700000/scsi@2/disk c) /pci@1e,600000/ide@d/cdrom d) /pci@1e,600000/ide@d/disk e) /pci@1d,700000/emlx@2,1/fp@0,0/disk f) /pci@1d,700000/emlx@2/fp@0,0/disk q) NO SELECTION Enter Selection, q to quit:</pre>
4	Create an alias for the bootable LUN. Enter: nvalias <i><alias_name></i> <i>/<boot_LUN_pathname></i> Example: This example creates an alias called “sanbootdisk”. <pre>{0} ok nvalias sanbootdisk /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b {0} ok nvstore {0} ok devalias sanbootdisk /pci@8,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b</pre>

Step	Action
5	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: Suppose 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>
6	Boot the LUN. Enter: <pre>boot <alias_name></pre> Example: The following command boots the alias for <code>/pci@1d,700000/emlx@2/fp@0,0/disk@w500a098786f7ced2,b</code> <pre>{0} ok boot sanbootdisk</pre>
7	Enable auto-boot. Enter: <pre>setenv auto-boot? true</pre>
8	Once you have booted the host on the newly created SAN LUN, you should enable the multipathing that is appropriate for your system. See the vendor documentation for your multipathing solution for more information.

Modify OpenBoot for the Emulex LPFC driver

Modifying OpenBoot for LPFC involves

- ◆ Running `sanlun` to document the information about your LUN.
- ◆ Mapping the HBA instance to a physical path.
- ◆ Binding the target and the LUN to the HBA.

- ◆ Creating an alias for the bootable LUN. This alias substitutes for the device address (/pci@1d,700000/lpfc@2/sd@0,f5 in the example that follows) during subsequent boot operations.

The examples that follow use this LUN information:

```
c5t0d245 <NETAPP-LUN-0.2 cyl 6142 alt 2 hd 16 sec 256>
      /pci@1d,700000/lpfc@2/sd@0,f5
```

To modify OpenBoot, complete the following steps.

Step	Action
1	Run the <code>san lun lun show -p</code> command and write down the information for the SAN boot LUN. Make sure you write down the correct device name so that it matches the name you used when you made changes to the <code>/etc/vfstab</code> on the SAN boot LUN. Note_____ The <code>san lun</code> command is in the <code>/opt/NTAP/SANToolkit/bin/</code> directory. Example: The <code>san lun lun show -p</code> command produces the following information. <pre># /opt/NTAP/SANToolkit/bin/san lun lun show -p fas3020-shu01:/vol/vol245/lun245 (LUN -1) 12g (12884901888) LUN State: GOOD Filer_CF_State: Cluster Enabled Filer Partner: fas3020-shu02 Multipath_Policy: None Multipath-provider: None ----- host filer path path device host primary partner state type filename HBA filer filer ----- up secondary /dev/rdisk/c6t3d245s2 lpfc1 0d up secondary /dev/rdisk/c6t2d245s2 lpfc1 0c up primary /dev/rdisk/c6t1d245s2 lpfc1 0d up primary /dev/rdisk/c6t0d245s2 lpfc1 0c up secondary /dev/rdisk/c5t3d245s2 lpfc0 0b up secondary /dev/rdisk/c5t2d245s2 lpfc0 0a up primary /dev/rdisk/c5t1d245s2 lpfc0 0b up primary /dev/rdisk/c5t0d245s2 lpfc0 0a</pre>

Step	Action
2	Get the device path for the lpfc instance to device path mapping from the /etc/path_to_inst file. Example: You can use the grep command to get this information. <pre># grep \"lpfc\" /etc/path_to_inst</pre> <pre>"/pci@1d,700000/lpfc@2" 0 "lpfc" "/pci@1d,700000/lpfc@2,1" 1 "lpfc"</pre>
3	Document the WWPN binding for your controllers. The persistent binding information can be found in the file /kernel/drv/lpfc.conf. Look for the variable nameds fcp-bind-WWPN. You should also look for the lpfc instance and target combination to match the path used in the /etc/vfstab file on the SAN boot LUN. Example: The binding for the sample configuration is provided below: <pre>fcp-bind-wwpn="500a098486c7ced2:lpfc0t0", "500a098386c7ced2:lpfc0t1", "500a098286c7ced2:lpfc1t0", "500a098186c7ced2:lpfc1t1", "500a098496c7ced2:lpfc0t2", "500a098396c7ced2:lpfc0t3", "500a098296c7ced2:lpfc1t2", "500a098196c7ced2:lpfc1t3";</pre>
4	Activate the OpenBoot environment by performing the following steps: 1. Halt the Solaris operating system by entering the command: halt Result: The Open Boot program displays the OpenBoot prompt. 2. Stop automatic reboot during OpenBoot configuration by entering the command: setenv auto-boot? false 3. Reset the host and activate your changes by entering: reset-all

Step	Action
5	Bind the HBA to the LUN and the target. Note This step uses all of the information gathered in steps 1 through 3. <pre>select /device_path_to_card wwpn <wwpn of target> <lun in hex> <target in hex> set-boot-id reset-all</pre> Example: The following example show the type of output that occurs. <pre>{1} ok select /pci@1d,700000/lpfc@2 {1} ok wwpn 500a098486c7ced2 f5 0 set-boot-id</pre> <pre>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 0 Lnk_timer f Plogi-timer 0 LUN (1 byte) f5 DID 0 WWPN 500a.0984.86c7.ced2 LUN (8 bytes) 0000.0000.0000.0000</pre> <pre>*** Type reset-all to update. *** {1} ok reset-all</pre>

Step	Action
6	Run the show-disks command to get the correct device path name for your Fibre Channel card. It is important to look for the device path associated with the LUN you used during the earlier format step. Due to the way the OBP sees hardware, the device may show up with either “disk” or “ssd” as the terminating characters and those characters may differ from what was seen in the format output. Example: This output shown here continues the previous example and uses the information that was written down during the format command. The LUN is “w500a098786f7ced2,b” and the device path is “/pci@1d,700000/emlx@2/fp@0,0/ssd.” <pre>4. c2t500A098786F7CED2d11 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b</pre> The following example simply shows the type of output you see when you execute the show-disks command. <pre>{1} ok show-disks a) /pci@1f,700000/scsi@2,1/disk b) /pci@1f,700000/scsi@2/disk c) /pci@1e,600000/ide@d/cdrom d) /pci@1e,600000/ide@d/disk e) /pci@1d,700000/emlx@2,1/fp@0,0/disk f) /pci@1d,700000/emlx@2/fp@0,0/disk q) NO SELECTION Enter Selection, q to quit:</pre>
7	Create an alias for the bootable LUN. Enter: nvalias <i><alias_name></i> <i>/<boot_LUN_pathname></i> Example: This example creates an alias called “sanbootdisk”. <pre>0} ok nvalias sanbootdisk /pci@1d,700000/lpfc@2/sd@0,f5 {0} ok nvstore {0} ok devalias sanbootdisk /pci@1d,700000/lpfc@2/sd@0,f5</pre>

Step	Action
8	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: Suppose 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>
9	Enable auto-boot. Enter: <pre>setenv auto-boot? true</pre>
10	Boot the LUN. Enter: <pre>boot <alias_name></pre> Example: The following command boots the alias for <code>/pci@1d,700000/lpfc@2/sd@0,f5</code> <pre>{0} ok boot sanbootdisk</pre>
11	Once you have booted the host on the newly created SAN LUN, you should enable the multipathing that is appropriate for your system. See the vendor documentation for your multipathing solution for more information.

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 136
- ◆ [“Completing prerequisites”](#) on page 138
- ◆ [“Upgrading to single image mode”](#) on page 139

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 FC hosts.
- ◆ There are no switch limitations. You can connect an NetApp storage system in single image mode to any supported FC 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.
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>sanlun lun show</code> - display LUNs visible to the host. ◆ <code>sanlun 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 <i>diskgroup_name</i></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 FC-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 FC 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>fcv 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 FC service on each system in the cluster. On each system, enter the following command: <pre>fcv 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 “(LPFC drivers) Creating persistent bindings” on page 55 for additional information about using lputil. ◆ To use create_bindings.pl, enter: create_binding.pl -d -n storage_system_name Refer to Appendix B, “Using the create_binding.pl Script,” on page 83 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

About this appendix This appendix provides tips and troubleshooting information to help you as you use the FC 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 146
- ◆ [“Using Diagnostic Utilities”](#) on page 152
- ◆ [“Collecting storage system information”](#) on page 153
- ◆ [“Collecting Solaris information”](#) on page 155
- ◆ [“Collecting switch information”](#) on page 156

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 NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do>. 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 FC Solaris Host Utilities 4.0 or FC Solaris Attach Kit installed, you must remove it. See “[Removing an FC Solaris Host Utilities or Attach Kit installation](#)” on page 47.

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.

```
# sanlun fcp show adapter -v

adapter name:      emlxs0
WWPN:              10000000c94be338
WWNN:              20000000c94be338
driver name:       emlxs
model:             LP11002-M4
model description: Emulex LP11002-M4 4Gb 2port FC: PCI-X2 SFF HBA
serial number:     VM53824441
hardware version:  1036406d
driver version:    2.20k (2007.06.04.09.35)
firmware version:  2.10a10
Number of ports:   1 of 2
port type:         Fabric
port state:        Operational
supported speed:   1 Gbit/sec, 2 Gbit/sec, 4 Gbit/sec
negotiated speed:  4 Gbit/sec
OS device name:    /dev/cfg/c0

adapter name:      emlxs1
WWPN:              10000000c94be339
WWNN:              20000000c94be339
driver name:       emlxs
model:             LP11002-M4
model description: Emulex LP11002-M4 4Gb 2port FC: PCI-X2 SFF HBA
serial number:     VM53824441
hardware version:  1036406d
driver version:    2.20k (2007.06.04.09.35)
firmware version:  2.10a10
Number of ports:   2 of 2
```

```

port type:          Fabric
port state:         Operational
supported speed:    1 Gbit/sec, 2 Gbit/sec, 4 Gbit/sec
negotiated speed:   4 Gbit/sec
OS device name:     /dev/cfg/c1

```

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   disk0         dg            online
FAS30200_1  auto:cdsdisk   disk1         dg            online
FAS30200_2  auto:cdsdisk   disk2         dg            online
FAS30200_3  auto:cdsdisk   disk3         dg            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 “[Prerequisites](#)” 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.
- ◆ 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.

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 153 for more information.
- ◆ `solaris_info`. See “[Collecting Solaris information](#)” on page 155 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 156 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 is 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 FC 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 /opt/NTAP/SANToolkit/bin 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 is 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
 setting variables 39

A

A/A
 Active/Active paths 76
A/P-C
 Active/Passive Concurrent paths 76
active paths
 Active/Active (A/A) 76
Active/Passive Concurrent paths
 A/P-C 76
Array Support Library
 see ASL 66
ASL
 components 66
 definition 66
 fresh install 68
 multipathing reports 76
Attach Kits
 previous kits 2

B

basic_config
 configuration script 3
 description 37
 example with LPFC drivers 41
 options (LPFC drivers) 40
 options (native drivers) 43
boot data
 copying 122
bootblk
 installing 122

C

cfmode
 effect of changing 136

 identifying 38
 preparing host for upgrade 139
 prerequisites for changing 138
 reconfiguring host 142
 single image requirements 136
 upgrading overview 139
 upgrading storage cluster 140
 validating zoning 141

checklists

 installation tasks 5

commands

 vxctl enable 63
 vxdiskadm 63
 vxdsk list 148
 vxprint 148

configuring system parameters 37

create_binding.pl

 definition 84
 prerequisites 86
 rules for using 84
 using 88

D

diagnostic utilities

 <switch>_info 152
 collecting Solaris information 155
 collecting storage system information 153
 collecting switch information 156
 executing 152
 filer_info 152
 solaris_info 152
 switch_info 156

disk shelves

 powering up 150

documentation

 obtaining 34

drivers

 downloading software (Emulex) 17
 Sun drivers for Emulex 23

E

Emulex

- changing HBA mode 105
- changing HBA to Emulex mode 111
- downloading HBA software 15
- driver software 14
- Sun native drivers 23

Emulex HBA

- downloading software 17
- firmware versions 24, 26

equipment

- powering up 150

F

FC Solaris Attach Kit

- removing 48

FC Solaris Host Utilities

- configuration overview 52
- contents 2
- defined 2
- documentation 34
- obtaining 32
- SAN Toolkit 3
- upgrading 46

FC Solaris Host Utilitiest

- removing 47

FCode

- enabling compatibility (QLogic) 109
- upgrading (Emulex) 24
- upgrading (QLogic) 27

filer_info

- diagnostic script 3

firmware

- upgrading 22
- version 24, 26

format utility

- using to label LUNs 61

H

HBA

- checking 147

host

- communicating with storage system 151

- powering up 150

I

igroup

- creating 54
- definition 52

L

LPFC drivers

- basic_config options 40
- changing HBA to SD mode 111
- configuring parameters 38
- discovering LUNs 59
- firmware versions 20
- modifying sd.conf 58

lpfc.conf file

- setting variables 38

lputil

- creating persistent bindings 56

LUNs

- creating 54
- creating bootable LUN 115
- discovering (LPFC drivers) 58, 59
- discovering (native drivers) 60
- labeling 61
- partitioning bootable LUN 115

M

mpxio_set

- configuration script 3

multipathing

- displaying information 96

N

native drivers

- configuring 43
- discovering LUNs 60
- Emulex 23
- getting controller number
HBAs
getting controller number 60
- setting variables 43

- upgrading FCode (Emulex) 24
- upgrading FCode (QLogic) 27

O

OpenBoot

- definition 126
- modifying 126

P

persistent binding

- creating 84
- creating with create_binding.pl 88
- creating with lputil 56
- definition 52, 55
- methods for creating 55

- powering up equipment 150

Q

QLogic

- changing HBA mode 109
- enabling FCode compatibility 109
- upgrading FCode 27

R

removing

- attach kit 48
- FC Solaris Host Utilities 47

- removing Host Utilities 47

S

SAN boot

- changing Emulex HBA mode 105
- changing HBA to SD mode(LPFC) 111
- changing QLogic HBA mode 109
- configuration overview 103
- copying the boot data 122
- creating bootable LUN 115
- definition 102
- installing the bootblk 122
- modifying OpenBoot 126
- partitioning bootable LUN 115

- performing direct install 121

- prerequisites 103

SAN Toolkit 3

- components 3
- san_version command 3
- sanlun utility 3

san_version

- command 3

sanlun

- part of SAN Toolkit 3
- syntax 97
- using 93

scripts

- basic_config 3
- filer_info 3
- mpxio_set 3

sd.conf file

- configuring to discover LUNs 58
- definition 52

single image mode

- changing 136
- configuration overview 139
- definition 136
- preparing the host 139
- prerequisites 138
- reconfiguring the host 142
- Solaris requirements 136
- upgrading the storage system 140
- validating zoning 141

Solaris

- displaying HBA information 98
- displaying LUN information 94
- preparing for upgrade to single image mode 139
- requirements for single image mode 136

ssd.conf file

- setting variables 43

storage system

- communicating with host 151
- powering up 150

switch_info

- executing 156

switches

- powering up 150
- verifying setup 150

- zoning 150
- system parameters
 - configuring 37
 - configuring manually 38
 - pre-requisites 36

U

- uninstall script 47, 48
- upgrading
 - Host Utilities 46

V

- Veritas Volume Manager
 - check VxVM 148
 - configuring LUNs 63

Z

- zoning 150
 - validating for single image mode 141