

Solaris Host Utilities 5.1

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: 215-04881_A0
August 2009

Contents

Copyright information.....	11
Trademark information.....	13
About this guide.....	15
Audience.....	15
Terminology.....	15
Where to enter commands.....	16
Command, keyboard, and typographic conventions.....	16
Special messages.....	17
How to send your comments.....	17
Changes to this document: August 2009.....	19
The Solaris Host Utilities.....	21
Overview of the Solaris Host Utilities.....	21
Introduction to the Solaris Host Utilities.....	21
Supported Solaris environments and protocols.....	21
How to find instructions for your Solaris environment.....	23
Solaris Host Utilities general components.....	24
What the Host Utilities contain.....	24
Protocols and configurations supported by the Solaris Host Utilities.....	26
Notes about the supported protocols.....	26
The FC protocol.....	26
The iSCSI protocol.....	26
Supported configurations.....	27
Supported Solaris and Data ONTAP features.....	27
Features supported by the Host Utilities.....	28
HBAs and the Solaris Host Utilities.....	28
Multipathing and the Solaris Host Utilities.....	29
iSCSI and multipathing.....	29
Volume managers and the Solaris Host Utilities.....	29
(MPxIO/FC) ALUA support with certain versions of Data ONTAP.....	30
Sun Microsystems' Logical Domains and the Host Utilities.....	30
SAN booting and the Host Utilities.....	30
Support for non-English versions of Solaris operating systems.....	31

High-level look at Host Utilities' Veritas DMP stack.....	31
High-level look at Host Utilities' MPxIO stack.....	32
Where to find more information.....	33
Planning the installation and configuration of the Host Utilities.....	37
Prerequisites for installing and setting up the Solaris Host Utilities.....	37
Installation overview for the Solaris Host Utilities.....	38
iSCSI configuration overview.....	39
LUN configuration overview.....	39
(Veritas DMP/FC) Information on setting up the drivers.....	41
General information on getting the driver software.....	41
Emulex LPFC drivers.....	42
Emulex components for LPFC drivers.....	42
Downloading and extracting the Emulex software.....	42
Special Note about the Emulex LPFC driver utility hbacmd.....	43
Information on working with drivers prior to 6.21g.....	44
Information on working with drivers 6.21g.....	47
Determining the Emulex LPFC firmware and FCode versions.....	49
Upgrading the LPFC firmware and FCode.....	49
Sun drivers for Emulex HBAs (emlxs).....	50
Installing the EMLXemlxu utilities.....	50
Determining the Emulex firmware and FCode versions for native drivers.....	51
Upgrading the firmware for native drivers.....	51
Updating your FCode HBAs with native drivers.....	52
Sun drivers for QLogic HBAs (qlc).....	52
Downloading and extracting the QLogic software.....	52
Installing the SANsurfer CLI package.....	53
Determining the FCode on QLogic cards.....	53
Upgrading the QLogic FCode.....	54
The Solaris Host Utilities installation process.....	57
Key steps involved in setting up the Host Utilities.....	57
The software packages.....	58
Downloading the Host Utilities software from the NOW site.....	58
Installing the Solaris Host Utilities software.....	59
Information on upgrading or removing the Solaris Host Utilities.....	63
Upgrading the Solaris Host Utilities or reverting to another version.....	63

Methods for removing the Solaris Host Utilities.....	63
Uninstalling Solaris Host Utilities 5.x, 4.x, 3.x.....	64
Uninstalling the Attach Kit 2.0 software.....	65
(iSCSI) Additional configuration for iSCSI environments.....	67
iSCSI node names.....	67
(iSCSI) Recording the initiator node name.....	68
(iSCSI) Storage system IP address and iSCSI static, ISNS, and dynamic discovery	68
(MPxIO/iSCSI) Configuring MPxIO with mpxio_set in iSCSI environments.....	69
(Veritas DMP/iSCSI) Support for iSCSI in a Veritas DMP environment.....	69
(Veritas DMP/iSCSI) Disabling MPxIO when using iSCSI with Veritas DMP.....	70
(iSCSI) CHAP authentication.....	71
(iSCSI) Configuring bidirectional CHAP.....	71
(iSCSI) Data ONTAP upgrades may affect CHAP configuration.....	72
(Veritas DMP/FC) Tasks for completing the setup of a Veritas DMP stack using FC.....	73
(Veritas DMP/FC) Before you configure the Host Utilities for Veritas DMP.....	73
(Veritas DMP/FC) About the basic_config command for Veritas DMP environments.....	74
(Veritas DMP/LPFC) Tasks involved in configuring systems using LPFC drivers.....	75
(Veritas DMP/LPFC) Identifying the cfmode.....	75
(Veritas DMP/LPFC) Values for the lpfc.conf variables.....	76
(Veritas DMP/LPFC) Values for the /etc/system variables on systems using LPFC.....	76
(Veritas DMP/LPFC) The basic_config options on systems using LPFC drivers.....	77
(Veritas DMP/LPFC) Example: Using basic_config on systems with LPFC drivers.....	78
(Veritas DMP/native) Tasks involved in configuring systems using native drivers.....	79
(Veritas DMP/native) ssd.conf variables for systems using native drivers.....	80

(Veritas DMP/native) About the basic_config command for native drivers.....	80
(Veritas DMP/native) The basic_config options on systems using native drivers and Veritas 5.0.....	81
(Veritas DMP/native) Example: Using basic_config on systems with native drivers and DMP.....	82
(MPxIO/FC) Tasks for completing the setup of a MPxIO stack.....	83
(MPxIO/FC) Before configuring system parameters on a MPxIO stack using FC.....	83
(MPxIO/FC) Preparing for ALUA for MPxIO in FC environments.....	84
(MPxIO/FC) Parameter values for systems using MPxIO with FC.....	85
(MPxIO/FC) About the basic_config command for MPxIO environments using FC.....	85
(MPxIO/FC) The basic_config options on systems using MPxIO with FC.....	86
Example: Executing basic_config on systems using MPxIO with FC.....	87
(Veritas DMP) Array Support Library and Array Policy Module.....	89
(Veritas DMP) About the Array Support Library and the Array Policy Module.....	89
(Veritas DMP) Information provided by the ASL.....	90
(Veritas DMP) Information on upgrading the ASL and APM	90
(Veritas DMP) ASL and APM installation overview.....	91
(Veritas DMP) Determining the ASL version.....	91
(Veritas DMP) How to obtain the ASL and APM	92
(Veritas DMP) Installing the ASL and APM software.....	92
(Veritas DMP) Tasks to perform before you uninstall the ASL and APM.....	94
(Veritas DMP) What an ASL array type is.....	98
(Veritas DMP) The storage system's FC failover mode or iSCSI configuration and the array types.....	99
(Veritas DMP) How sanlun displays the array type.....	99
(Veritas DMP) Using VxVM to display available paths.....	99
(Veritas DMP) Using sanlun to obtain multipathing information.....	101
(Veritas DMP) Veritas environments and the fast recovery feature	102
(Veritas DMP) The Veritas DMP restore daemon requirements.....	102
(Veritas DMP) Setting the restore daemon interval.....	102

(Veritas DMP) Information on ASL error messages	103
LUN configuration and the Solaris Host Utilities.....	105
Overview of LUN configuration and management.....	105
Tasks necessary for creating and mapping LUNs.....	107
How the LUN type affects performance.....	108
Methods for creating igroups and LUNs.....	108
Best practices for creating igroups and LUNs.....	109
(Veritas DMP) LPFC drivers and LUNs.....	109
(Veritas DMP/LPFC) Persistent bindings.....	109
(Veritas DMP/LPFC) Methods for creating persistent bindings.....	110
(Veritas DMP/LPFC) Creating WWPN bindings with hbacmd.....	110
(LPFC drivers prior to 6.21g) Modifying the sd.conf file	111
(LPFC drivers prior to 6.21g) Discovering LUNs.....	111
(LPFC drivers 6.21g) Discovering LUNs.....	112
(iSCSI) Discovering LUNs.....	113
Sun native drivers and LUNs.....	113
(Sun native drivers) Getting the controller number.....	114
(Sun native drivers) Discovering LUNs.....	114
Labeling the new LUN on a Solaris host.....	115
Methods for configuring volume management.....	117
The sanlun utility.....	119
Displaying host LUN information with sanlun.....	119
Displaying path information with sanlun.....	121
Explanation of the sanlun lun show -p output.....	122
Displaying host HBA information with sanlun.....	123
Troubleshooting.....	127
Check the Release Notes.....	127
About the troubleshooting sections that follow.....	128
Check the version of your host operating system.....	129
Confirm the HBA is supported.....	129
(MPxIO, native drivers) Ensure that MPxIO is configured correctly for ALUA on FC systems.....	131
(MPxIO, FC) Ensure that MPxIO is enabled on SPARC systems.....	131
(MPxIO) Ensure that MPxIO is enabled on iSCSI systems.....	132
(MPxIO) Verify that MPxIO multipathing is working.....	133
(Veritas DMP) Check that the ASL and APM have been installed.....	134

(Veritas) Check VxVM.....	134
(MPxIO) Check the Solaris Volume Manager.....	136
(MPxIO) Check settings in ssd.conf or sd.conf.....	136
Check the storage system setup.....	137
(MPxIO/FC) Check the ALUA settings on the storage system.....	137
Verify that the switch is installed and configured.....	138
Determine whether to use switch zoning.....	138
Power up equipment in the correct order.....	139
Verify that the host and storage system can communicate.....	139
Possible iSCSI issues.....	139
(iSCSI) Verify the type of discovery being used.....	139
(iSCSI) Bidirectional CHAP doesn't work.....	140
(iSCSI) LUNs are not visible on the host.....	140
Diagnostic utilities for the Host Utilities.....	141
Executing the diagnostic utilities (general steps).....	142
Collecting host diagnostic information.....	142
Collecting storage system diagnostic information.....	143
Collecting switch diagnostic information.....	144
Possible MPxIO issues.....	145
(MPxIO) sanlun does not show all adapters.....	146
(MPxIO) Solaris log message says data not standards compliant.....	146
(Veritas DMP/LPFC) The Solaris Host Utilities	
create_binding program.....	147
(Veritas DMP/LPFC) The create_binding program.....	147
The rules create_binding uses for WWPN bindings.....	148
(Veritas DMP/LPFC) Prerequisites for using create_binding.....	148
(Veritas DMP/LPFC) Storage system prerequisites for using create_binding.....	149
(Veritas DMP/LPFC) Creating WWPN bindings with create_binding.....	149
SAN boot LUNs in a Solaris Veritas DMP environment	
with FC.....	153
(Veritas DMP) Prerequisites for creating a SAN boot LUN	154
(Veritas DMP) SAN boot configuration overview.....	154
(Veritas DMP/native) About setting up the Sun native HBA for	
SAN booting.....	155
(Veritas DMP/native) Changing the Emulex HBA to SFS mode.....	155

(Veritas DMP/native) Changing the QLogic HBA to enable FCode compatibility.....	158
(Veritas DMP/LPFC) About setting up the LPFC HBA for SAN booting.....	160
(Veritas DMP/LPFC) Changing the Emulex HBA to SD mode.....	160
(Veritas DMP/native) Methods for installing directly onto a SAN	
boot LUN.....	162
(Veritas DMP) Installing the bootblk.....	163
(Veritas DMP) Copying the boot data.....	163
(Veritas DMP) Information on creating the bootable LUN.....	165
(Veritas DMP) Partitioning the bootable LUN to match the host device.....	166
(Veritas DMP) What OpenBoot is	169
(Veritas DMP/native) Modifying OpenBoot with Sun native drivers.....	170
(Veritas DMP/LPFC) Modifying OpenBoot for the Emulex LPFC driver.....	172
SAN boot LUNs in an MPxIO environment.....	177
(MpxIO) Prerequisites for creating a SAN boot LUN	178
(MPxIO) Options for setting up SAN booting.....	178
(MPxIO) Performing a direct install to create a SAN boot LUN.....	178
SPARC systems without MPxIO: Copying data from locally booted server.....	179
x86/x64 with MPxIO systems: Copying data from a locally booted disk.....	182
x86/x64 without MPxIO systems: Copying data from locally booted server.....	185
Index.....	189

Copyright information

Copyright © 1994–2009 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

All applicable trademark attribution is listed here.

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; Shadow Tape; 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, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>.

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.

NetCache is certified RealSystem compatible.

About this guide

Here you can learn what this document describes and who it is intended for, what special terminology is used in the document, what command, keyboard, and typographic conventions this document uses to convey information, and other details about finding and using information.

This guide describes how to configure a Solaris® host in a SAN environment. The host accesses data on a storage system that runs Data ONTAP® software with the Fibre Channel (FC) or iSCSI protocol.

See the NetApp Interoperability Matrix for information on the specific Solaris versions supported.

Next topics

[Audience](#) on page 15

[Terminology](#) on page 15

[Where to enter commands](#) on page 16

[Command, keyboard, and typographic conventions](#) on page 16

[Special messages](#) on page 17

[How to send your comments](#) on page 17

Related information

[NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability/](http://now.netapp.com/NOW/products/interoperability/)

Audience

This guide is intended for system administrators who are familiar with Solaris.

You should be familiar with the specifics of your configuration, including the following items:

- Are familiar with the Sun® Solaris operating system
- Are familiar with either the MPxIO environment or the Veritas Storage System environment
- 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 document uses the following terms.

- Storage systems that run Data ONTAP are sometimes referred to as *filers*, *appliances*, *storage appliances*, or *systems*. The name of the FilerView graphical user interface for Data ONTAP reflects one of these common usages.
- LUN (logical unit) refers to a logical unit of storage.
- LUN ID refers to the numerical identifier for a LUN.

Where to enter commands

You can use your product more effectively when you understand how this document uses command conventions to present information.

You can perform common administrator tasks in one or more of the following ways:

- You can enter commands either at the system console or from any client computer that can obtain access to the storage system using a Telnet or Secure Shell (SSH) session.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.
- You can enter Windows, ESX, HP-UX, AIX, Linux, and Solaris commands at the applicable client console.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.

Command, keyboard, and typographic conventions

This document uses command, keyboard, and typographic conventions that help you enter commands.

Command conventions

In examples that illustrate commands executed on a UNIX workstation, the command syntax and output might differ, depending on your version of UNIX.

Keyboard conventions

- When describing key combinations, this document uses the hyphen (-) to separate individual keys. For example, "Ctrl-D" means pressing the "Control" and "D" keys simultaneously.
- This document uses the term "Enter" to refer to the key that generates a carriage return, although the key is named "Return" on some keyboards.

Typographic conventions

The following table describes typographic conventions used in this document.

Convention	Type of information
<i>Italic font</i>	Words or characters that require special attention. Placeholders for information you must supply. For example, if the guide says to enter the <code>arp -d <i>hostname</i></code> command, you enter the characters "arp -d" followed by the actual name of the host. Book titles in cross-references.
Monospaced font	Command names, option names, keywords, 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 you must type it in uppercase letters.

Special messages

This document might contain the following types of messages to alert you to conditions you need to be aware of. Danger notices and caution notices only appear in hardware documentation, where applicable.

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

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

Danger: A danger notice warns you of conditions or procedures that can result in death or severe personal injury.

Caution: A caution notice warns you of conditions or procedures that can cause personal injury that is neither lethal nor extremely hazardous.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by e-mail to docomments@netapp.com. To help us direct your comments to the correct division, include in the

subject line the name of your product and the applicable operating system. For example, *FAS6070—Data ONTAP 8.0*, or *Host Utilities—Solaris*, or *Operations Manager 3.8—Windows*.

Changes to this document: August 2009

Several changes have been made to this document since it was published for the Solaris Host Utilities 5.0 release.

This document has been updated to add the following information:

- iSCSI and Fibre Channel (FC) information has been combined in this manual. Previously, the Host Utilities provided separate installation and setup guides for people using the iSCSI protocol and for people using the FC protocol. This manual puts the information from both those manuals into one installation and setup guide.
- iSCSI is supported with certain Veritas Dynamic Multipathing (DMP) environments. You cannot use iSCSI with Veritas DMP if MPxIO is enabled. If you previously installed a version of the Host Utilities that used MPxIO, you may need to remove certain settings from the `/kernel/drv/scsi_vhci.conf` file before you can use DMP.
- `controller_info` has been added as a diagnostic utility. `controller_info`, like `filer_info`, gathers information about the storage system. However, `controller_info` uses the ZAPI interface while `filer_info` uses the Remote Shell (RSH). Data ONTAP 8.0 and later use the ZAPI interface as the default so `controller_info` will work without requiring additional setup steps.
- The `create_binding` program has been updated to use the ZAPI interface.
- Several Host Utilities scripts have been compiled. They include:
 - `create_binding`
 - `filer_info`
 - `brocade_info`
 - `cisco_info`
 - `mcdata_info`
 - `qlogic_info`
 - `controller_info`

This change means that you need to enter `create_binding` instead of `create_binding.pl`. There are no changes in how you execute the other diagnostic utilities.

- The Solaris Host Utilities support non-English versions of the Solaris operating system.

The Solaris Host Utilities

The Solaris Host Utilities are a set of tools you can use to connect a Solaris host to a storage system.

Next topics

[Overview of the Solaris Host Utilities](#) on page 21

[Protocols and configurations supported by the Solaris Host Utilities](#) on page 26

[Supported Solaris and Data ONTAP features](#) on page 27

[Where to find more information](#) on page 33

Overview of the Solaris Host Utilities

The tools provided by the Host Utilities support several Solaris environments.

Next topics

[Introduction to the Solaris Host Utilities](#) on page 21

[Supported Solaris environments and protocols](#) on page 21

[How to find instructions for your Solaris environment](#) on page 23

[Solaris Host Utilities general components](#) on page 24

[What the Host Utilities contain](#) on page 24

Introduction to the Solaris Host Utilities

The Solaris Host Utilities are a collection of components that enable you to connect Solaris hosts to NetApp storage systems running Data ONTAP. Once connected you can set up logical storage units (LUNs) on the storage system.

Note: Previous versions of the Host Utilities were called FCP Solaris Attach Kits and iSCSI Support Kits.

The following sections provide an overview of the Solaris Host Utilities as well as information on what components the Host Utilities supply.

Supported Solaris environments and protocols

The Host Utilities support several Solaris environments.

For details on which environments are supported, see the online NetApp Interoperability Matrix.

The following table summarizes key aspects of the two main environments.

Solaris Environment	Notes
Veritas DMP	• This environment uses Veritas Storage Foundation and its features. • Multipathing: Veritas Dynamic Multipathing (DMP) with either Emulex® LightPulse® Fibre Channel (LPFC) driver or Sun native drivers (Leadville) or iSCSI. • Volume management: Veritas Volume Manager (VxVM). • Protocols: Fibre Channel (FC) and iSCSI. • Software package: Install the software packages in the compressed download file for your processor. Note: At the time this document was produced, the Host Utilities only supported the x86/64 processor with Veritas when you used the iSCSI protocol. It did not support it with the FC protocol. • Setup issues: • You may need to perform some driver setup. • The Symantec Array Support Library (ASL) and Array Policy Module (APM) must be installed. • If you are using iSCSI and have used a previous version of the Host Utilities, you must make sure MPxIO is disabled. • Configuration issues: • Systems using LPFC drivers require changes to the parameters in the <code>/etc/system</code>, <code>/kernel/drv/lpfc.conf</code>, and <code>/kernel/drv/sd.conf</code> files. • Systems using Sun native FC drivers require changes to the parameters in the <code>/kernel/drv/ssd.conf</code> file. Note: At the time this document was produced, the Veritas environment did not support Asymmetric Logical Unit Access (ALUA).

Solaris Environment	Notes
MPxIO (Native OS)	- This environment works with features provided by the Solaris operating system. It uses Sun StorEdge SAN Foundation Software. - Multipathing: Sun StorageTek Traffic Manager (MPxIO) or the Solaris iSCSI Software Initiator. - Volume management: Sun Volume Manager (SVM), ZFS, or VxVM. - Protocols: FC and iSCSI. - ALUA is only supported in FC environments. (It is also supported with one older version of the iSCSI Support Kit: 3.0.) - Software package: Download the compressed file associated with your system's processor (SPARC or x86/64) and install the software packages in that file. - Setup issues: None. - Configuration issues: - Systems using SPARC processors and FC require changes to the parameters in the <code>/kernel/drv/ssd.conf</code> file.

Related information

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

How to find instructions for your Solaris environment

Many of the instructions in this manual apply to all the Solaris environments that the Host Utilities support. In some cases, the commands you use vary based on the environment you are using.

If information or instructions for a step or feature apply only to one environment, this guide notes that fact. To make it easy to quickly identify information for your environment, when a section applies to a specific Host Utilities environment, this guide identifies the environment by placing a qualifier in the headings. This guide uses the following qualifiers:

Qualifier	The section that follows applies to
(Veritas DMP)	Environments using Veritas DMP as the multipathing solution.
(Veritas DMP/LPFC)	Veritas DMP environments that use LPFC drivers.
(Veritas DMP/native)	Veritas DMP environments that use Sun native drivers.
(Veritas DMP/FC)	Veritas DMP environments that use the FC protocol.
(Veritas DMP/iSCSI)	Veritas DMP environments that use the iSCSI protocol.
(MPxIO)	Environments using MPxIO as the multipathing solution. Currently, all MPxIO environments use native drivers.

Qualifier	The section that follows applies to
(MPxIO/FC)	MPxIO environments using the FC protocol.
(MPxIO/iSCSI)	MPxIO environments using the iSCSI protocol.
(iSCSI)	Environments using the iSCSI protocol.

There is also information about using the Host Utilities in a Solaris environment in the *Release Notes* and Solaris Host Utilities reference documentation. For information on the documents associated with the Host Utilities, see the SAN/IPSAN Information Library page on the NOW site. You can download all the Host Utilities documentation from the NOW site.

Related information

[SAN/IPSAN Information Library page - http://now.netapp.com/NOW/knowledge/docs/san/](http://now.netapp.com/NOW/knowledge/docs/san/)

Solaris Host Utilities general components

The Host Utilities consist of tools, diagnostic scripts, and documentation.

It consists of both software and documentation.

When you install the Host Utilities, it installs the SAN Toolkit, which contains:

- Configuration tools that help you configure host, HBA, system files, and create persistent bindings. You can configure the Solaris host using the tools provided in the Solaris Host Utilities, HBA-provided tools, or a combination Solaris Host Utilities and HBA-provided tools.
- The `sanlun` utility, which helps you manage LUNs and the host HBA.

The documentation that comes with the Host Utilities includes:

- This installation and setup guide
- *Release Notes*
- *Host Settings Changed by the Host Utilities*
- *Quick Command Reference*

Note: The *Release Notes* are updated whenever new information relating to the Host Utilities is found. You should check the *Release Notes* before installing the Host Utilities to see if there are new details or corrections that apply to the installation and periodically after that to see if new information about the Host Utilities has been discovered.

What the Host Utilities contain

The Host Utilities bundle the software tools into a SAN Toolkit.

Note: This toolkit is common across all the Solaris Host Utilities configurations: MPxIO and Veritas DMP as well as FC and iSCSI. As a result, some of its contents apply to one configuration, but not

another. Having a program or file that does not apply to your configuration does not affect performance.

The toolkit contains the following components:

- **basic_config** command line tool. This tool modifies the FC tunable variables to set them to values required by the FC driver. It provides the HBA time-out settings for both Sun and Veritas LPFC driver configurations.
- **controller_info** utility. Like the **filer_info** utility, this utility collects information about the storage system. It uses the ZAPI interface instead of the Remote Shell (RSH). This tool works with Data ONTAP 7.2 and higher.
- **create_binding** program. 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** utility. Like the **controller_info** utility, this utility collects information about the storage system. It uses RSH instead of the ZAPI interface. Some versions of Data ONTAP, such as 8.0, use the ZAPI interface as the default. With those versions, use the **controller_info** utility to avoid having to do additional setup to enable the **filer_info** utility.
- **mpxio_setscript**. This script adds or deletes the symmetric-option and NetApp VID/PID in the `/kernel/drv/scsi_vhci.conf` file.
(iSCSI) If you are using the iSCSI protocol, use this script to configure the host for symmetrical access.
(FC) If you are using the FC protocol, use this script if you previously installed an iSCSI version of the Solaris Host Utilities or an Solaris Support kit. In that case, you use this script to remove the symmetric-option from the `scsi_vhci.conf` file.
- **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. The **basic_config** command calls this script to make setting changes for the FC and SCSI drivers. You should never run this script.
- **solaris_info** utility. This script collects configuration information about your host.
- Additional diagnostic utilities. The following utilities 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.
- The man pages for **sanlun** and the diagnostic utilities.

Protocols and configurations supported by the Solaris Host Utilities

The Solaris Host Utilities provide support for Fibre Channel and iSCSI connections to the storage system using direct-attached, fabric-attached, and network configurations.

Next topics

[Notes about the supported protocols](#) on page 26

[The FC protocol](#) on page 26

[The iSCSI protocol](#) on page 26

[Supported configurations](#) on page 27

Notes about the supported protocols

The FC and iSCSI protocols enable the host to access data on storage systems.

The storage systems are targets that have storage target devices called LUNs. The protocol enables the host to access the LUNs to store and retrieve data.

For more information about using the protocols with your storage system, see the *Data ONTAP Block Access Management Guide for iSCSI and FC* for your version of Data ONTAP.

The FC protocol

The FC protocol requires one or more supported host bus adapters (HBAs) in the host. Each HBA port is an initiator that uses FC to access the LUNs on the storage system. The HBA port is identified by a worldwide port name (WWPN).

You need to make a note of the WWPN so that you can supply it when you create an initiator group (igroup). To enable the host to access the LUNs on the storage system using the FC protocol, you must create an initiator group on the storage system and give it the WWPN as an identifier. Then, when you create the LUN, you map it to that initiator group. This mapping enables the host to access that specific LUN.

The iSCSI protocol

The iSCSI protocol is implemented on both the host and the storage system.

On the host, the iSCSI protocol is implemented over either the host's standard Ethernet interfaces or an HBA.

On the storage system, the iSCSI protocol can be implemented over the storage system's standard Ethernet interface using one of the following:

- A software driver that is integrated into Data ONTAP.

- (Data ONTAP 7.1 and later) An iSCSI target HBA or an iSCSI TCP/IP offload engine (TOE) adapter. You do not have to have a hardware HBA.

The connection between the initiator and target uses a standard TCP/IP network. The storage system listens for iSCSI connections on TCP port 3260.

Supported configurations

The Host Utilities support fabric-attached, direct-attached, and network-attached configurations.

The *iSCSI and Fibre Channel Configuration Guide* provides detailed information, including diagrams, about the supported topologies. There is also configuration information in the *Data ONTAP Block Access Management Guide for iSCSI and FC* for your version of Data ONTAP. Refer to those documents for complete information on configurations and topologies.

The basic configurations supported by the Host Utilities are:

- Fabric-attached storage area network (SAN). The Host Utilities support two variations of fabric-attached SANs:
 - A single-host FC connection from the HBA to the storage system through a single switch. A host is cabled to a single FC switch that is connected by cable to redundant FC ports on an active/active storage system configuration. A fabric-attached single-path host has one HBA.
 - Two or more FC connections from the HBA to the storage system through dual switches or a zoned switch. In this configuration the host has at least one dual-port HBA or two single-port HBAs. The redundant configuration avoids the single point of failure of a single-switch configuration. This configuration requires that multipathing be enabled.

Note: Use redundant configurations with two FC switches for high availability in production environments. However, direct FC connections and switched configurations using a single zoned switch might be appropriate for less critical business applications.

- FC direct-attached. A single host with a direct FC connection from the HBA to stand-alone or active/active storage systems.
- iSCSI direct-attached. One or more hosts with a direct iSCSI connection to stand-alone or active/active storage systems. The number of hosts that can be directly connected to a storage system or a pair of storage systems depends on the number of available Ethernet ports.
- iSCSI network-attached. In an iSCSI environment, all methods of connecting Ethernet switches to a network that have been approved by the switch vendor are supported. Ethernet switch counts are not a limitation in Ethernet iSCSI topologies. Refer to the Ethernet switch vendor documentation for specific recommendations and best practices.

Supported Solaris and Data ONTAP features

The Host Utilities work with both Solaris and Data ONTAP features.

Next topics

- Features supported by the Host Utilities* on page 28
- HBAs and the Solaris Host Utilities* on page 28
- Multipathing and the Solaris Host Utilities* on page 29
- iSCSI and multipathing* on page 29
- Volume managers and the Solaris Host Utilities* on page 29
- (MPxIO/FC) ALUA support with certain versions of Data ONTAP* on page 30
- Sun Microsystems' Logical Domains and the Host Utilities* on page 30
- SAN booting and the Host Utilities* on page 30
- Support for non-English versions of Solaris operating systems* on page 31
- High-level look at Host Utilities' Veritas DMP stack* on page 31
- High-level look at Host Utilities' MPxIO stack* on page 32

Features supported by the Host Utilities

The Host Utilities support a number of features and configurations available with Solaris hosts and storage systems running Data ONTAP. Your specific environment affects what the Host Utilities support.

Some of the supported features include:

- Multiple paths to the storage system when a multipathing solution is installed
- Veritas VxVM, Sun Volume Manager, and ZFS file system
- **(MPxIO)** ALUA
- Sun Microsystems' Logical Domains
- SAN booting

For information on which features are supported with which configurations, see theNetApp Interoperability Matrix Tool.

Related information

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

HBAs and the Solaris Host Utilities

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 theNetApp Interoperability Matrix Tool.

Related information

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

Multipathing and the Solaris Host Utilities

The Solaris Host Utilities support different multipathing solutions based on your configuration.

Using multipathing allows you to configure multiple network paths between the host and storage system. If one path fails, FC traffic continues on the remaining paths. For a host to have multiple paths to a LUN, you must have multipathing enabled.

The Veritas environment of the Host Utilities uses Veritas DMP to provide multipathing.

The MPxIO environment of the Host Utilities uses Sun's native multipathing solution (MPxIO).

Note: The Host Utilities also support IP Multipathing (IPMP). You do not need to perform any specific NetApp configuration to enable IPMP.

iSCSI and multipathing

You can use iSCSI with either Veritas DMP or MPxIO.

You should have at least two Ethernet interfaces on the storage system enabled for iSCSI traffic. Having two interfaces enables you to take advantage of multipathing. Each iSCSI interface must be in a different iSCSI target portal group.

In Veritas DMP environments, you must also disable MPxIO before you can use iSCSI. You must use DMP for multipathing when you are using Veritas.

For more information about using multipathing with iSCSI, see *Using iSCSI Multipathing in the Solaris 10 Operating System*.

Related information

Using iSCSI Multipathing in the Solaris 10 Operating System - <http://www.sun.com/blueprints/1205/819-3730.pdf>

Volume managers and the Solaris Host Utilities

The Solaris Host Utilities support different volume management solutions based on your environment.

The Veritas DMP environment uses Veritas Volume Manager (VxVM).

The MPxIO stack works with Sun Volume Manager (SVM), ZFS, and VxVM to enable you to have different volume management solutions.

Note: To determine which versions of VxVM are supported with MPxIO, see the NetApp Interoperability Matrix Tool

Related information

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

(MPxIO/FC) ALUA support with certain versions of Data ONTAP

The MPxIO environment requires that you have ALUA enabled for active/active storage controllers (clustered storage systems) using FC and a version of Data ONTAP that supports ALUA.

Stand-alone storage controllers provide symmetric access to LUNs and do not use ALUA.

Note: ALUA is also known as Target Port Group Support (TPGS).

ALUA defines a standard set of SCSI commands for discovering path priorities to LUNs on FC and iSCSI SANs. When you have the host and storage controller configured to use ALUA, it automatically determines which target ports provide optimized (primary) and unoptimized (secondary) access to LUNs.

Note: The Host Utilities do not support ALUA with iSCSI, unless you are running the iSCSI Solaris Host Utilities 3.0.

Check your version of Data ONTAP to see if it supports ALUA and check the NetApp Interoperability Matrix Tool to see if the Host Utilities support that version of Data ONTAP. NetApp introduced ALUA support with Data ONTAP 7.2 and single-image cfmode.

Related information

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

Sun Microsystems' Logical Domains and the Host Utilities

Certain configurations of the Host Utilities MPxIO stack support Sun Microsystems Logical Domains (LDom).

The supported configurations include guests that are I/O Domains or guests that have iSCSI configured. You must install the Host Utilities if a guest is using NetApp storage.

If you are using LDom, you must configure your system with the Host Utilities settings. You can use Host Utilities `basic_config` command do this or you can configure the settings manually.

A Solaris host running LDom accesses and uses NetApp storage exactly the same way any other Solaris host does.

For information on which configurations support LDom, see the NetApp Interoperability Matrix Tool.

Related information

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

SAN booting and the Host Utilities

The Host Utilities support SAN booting in both the Veritas and MPxIO environments. SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for a Solaris host.

Configuring SAN boot allows you to do the following

- Use the SAN for your boot needs.
- Use the SAN for both storage and boot needs, thereby consolidating and centralizing storage.

The steps you need to perform to create a SAN boot LUN differ based on your Host Utilities environment. If you decide to set up SAN booting, make sure you use the instructions for your environment.

Support for non-English versions of Solaris operating systems

Solaris Host Utilities are supported on all language versions of Solaris. All product interfaces and messages are displayed in English; however, all options accept Unicode characters as input.

High-level look at Host Utilities' Veritas DMP stack

The Host Utilities' Veritas DMP stack works with Solaris hosts running Veritas Storage Foundation.

The following is a high-level summary of the supported Veritas DMP stack at the time this document was produced.

Note: Check the NetApp Interoperability Matrix Tool for details and current information about the supported stack.

- Operating system:
 - Solaris 9 and 10
- Processor:
 - SPARC processor systems
 - x86/64 processor systems

Note: At the time this document was produced, the Host Utilities only supported Veritas on x86/64 processors in iSCSI environments, not FC environments. To see the current information on which processors Veritas supports, see the NetApp Interoperability Matrix Tool .

- FC HBA
 - Emulex LPFC HBAs and their Sun-branded equivalents
 - Certain Sun OEM QLogic® HBAs and their Sun-branded equivalents
 - Certain Sun OEM Emulex® HBAs and their Sun Branded equivalents
- iSCSI software initiators
- Drivers
 - Emulex drivers (LPFC)
 - Sun-branded Emulex drivers (emlxs)
 - Sun-branded QLogic drivers (qlc)

- Multipathing
 - Veritas DMP

The Host Utilities Veritas DMP stack also supports the following:

- Volume manager
 - VxVM
- Clustering
 - Veritas Cluster Server (VCS)

Related information

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

High-level look at Host Utilities' MPxIO stack

The Host Utilities' MPxIO stack works with Solaris hosts running Sun StorEdge SAN Foundation Software and components that make up the native stack.

The following is a high-level summary of the supported MPxIO stack at the time this document was produced.

Note: Check the NetApp Interoperability Matrix Tool for details and current information about the supported stack.

- Operating system:
 - Solaris 10
- Processor:
 - SPARC processor systems
 - x86/64 processor systems
- HBA
 - Certain QLogic HBAs and their Sun-branded equivalents
 - Certain Emulex HBAs and their Sun-branded equivalents
- Drivers
 - Bundled Sun StorEdge SAN Foundation Software Emulex drivers (emlxs)
 - Bundled Sun StorEdge San Foundation Software QLogic drivers (qlc)
- Multipathing

- Sun StorageTek Traffic Manager (MPxIO)

The Host Utilities MPxIO stack also supports the following:

- Volume manager
 - SVM
 - VxVM

Note: To determine which versions of VxVM are supported with MPxIO, see the Interoperability Matrix.
- ZFS
- Clustering
 - Sun Clusters. This kit has been certified using the Sun Cluster Automated Test Environment (SCATE)
 - Veritas Cluster Server (VCS)

Related information

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

Where to find more information

For additional information about host and storage system requirements, supported configurations, best practices, your operating system, and troubleshooting, see the documents listed in the following table.

If you need more information about...	Go to...
Known issues, troubleshooting, operational considerations, and post-release developments	The latest Host Utilities <i>Release Notes</i> Note: The <i>Release Notes</i> are updated more frequently than the rest of the documentation. You should always check the release notes before installing the Host Utilities to see if there have been any changes to the installation or setup process since this document was prepared. You should check them periodically to see if there is new information on using the Host Utilities. A summary of what has been updated and when is on the Release Notes index page.
The latest supported configurations	• The Interoperability Matrix. • System Configuration Guide.

If you need more information about...	Go to...
A summary of some of the commands you might use with the Host Utilities	• The <i>FC Solaris Host Utilities Quick Command Reference</i> • The <i>iSCSI Solaris Host Utilities Quick Command Reference</i>
Supported SAN topologies	The <i>Fibre Channel and iSCSI Configuration Guide</i> for your version of Data ONTAP software
Changes to the host settings that are recommended by the Host Utilities	<i>Host Settings Affected by the Host Utilities</i>
Configuring the storage system and managing SAN storage on it	• Data ONTAP documentation Index • <i>Best Practices for Reliability: New System Installation</i> • <i>Data ONTAP® Software Setup Guide</i> • <i>Data ONTAP® Block Access Management Guide for iSCSI and FC</i> • <i>Data ONTAP® Release Notes</i> • <i>Command Reference</i> • <i>FilerView® Online Help</i>
Verifying compatibility of a storage system with environmental requirements	<i>Site Requirements</i>
Upgrading Data ONTAP	<i>Data ONTAP® Upgrade Guide</i>
Best practices/configuration issues	NetApp Knowledge Base
Migrating the cfmode, if necessary	Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations
The iSCSI protocol	The iSCSI protocol standard is defined by RFC 3720. Refer to www.ietf.org for more information.
Installing and configuring the HBA in your host	Your HBA vendor documentation
Your host operating system and using its features, such as SVM, ZFS, or MPxIO	Refer to your operating system documentation. You can download Sun manuals in PDF format from the Sun website.
Veritas Storage Foundation and its features	Refer to the Veritas documentation
Working with Emulex	Refer to the Emulex documentation.
Working with QLogic	Refer to the QLogic documentation.
General product information, including support information	The NetApp NOW site at now.netapp.com

Related information

NetApp Interoperability Matrix - <http://now.netapp.com/NOW/products/interoperability/>

Fibre Channel and iSCSI Configuration Guide -

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/

NetApp documentation - <http://now.netapp.com/NOW/knowledge/docs/docs.cgi>

NOW Knowledge Base and Technical Documentation -

<http://now.netapp.com/NOW/knowledge/docs/san/docs.shtml>

Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations -

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/SSICFMODE_1205.pdf

Emulex partner site (when this document was produced) -

<http://www.emulex.com/support/solaris/index.jsp>

QLogic partner site (when this document was produced) -

http://support.qlogic.com/support/drivers_software.asp

Sun documentation (when this document was produced) - <http://docs.sun.com/>

Planning the installation and configuration of the Host Utilities

Installing the 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 a high-level look at the tasks you need to perform to install and configure the Host Utilities to help you in planning your installation and configuration.

Detailed information on installing and configuring the Host Utilities is provided in the chapters that follow this overview.

Note: Occasionally there are known problems that can affect your system setup. Please check the *Solaris Host Utilities Release Notes* before you install the Host Utilities to make sure there are no issues affecting your system setup. The *Release Notes* are updated whenever an issue is found and may contain more current information than this manual.

Next topics

[Prerequisites for installing and setting up the Solaris Host Utilities](#) on page 37

[Installation overview for the Solaris Host Utilities](#) on page 38

[iSCSI configuration overview](#) on page 39

[LUN configuration overview](#) on page 39

Prerequisites for installing and setting up the Solaris Host Utilities

There are several tasks you should perform before you install the Host Utilities

These tasks include the following:

1. Set up your system:

- Host OS and appropriate updates
- HBAs or software initiators
- Drivers
- **Veritas environments only:** Veritas Storage Foundation, the Array Support Library (ASL) for the storage controllers, and, with Veritas Storage Foundation 5.0, Array Policy Module (APM)

Note: You must install the Veritas VxVM before you install the ASL and APM software. These components are available from the Symantec web site.

- Volume management and multipathing, if used

- Storage system with Data ONTAP installed
- **FC environments only:** Switches, if used.

Note: For information on supported topologies, see the *iSCSI and Fibre Channel Configuration Guide*, which is available online.

For the most current information on system requirements, see the NetApp Interoperability Matrix Tool

2. Verify that your storage system is:
 - Licensed correctly for either FC service or the iSCSI service and running it.
 - Using the recommended cfmode (single-image).
 - Configured to work with the target HBAs, as needed by your protocol.
 - Set up to work with the Solaris host and the initiator HBAs or software initiators, as needed by your protocol.
 - **MPxIO, FC environments only:** Set up to work with ALUA, if it is supported with your version of Data ONTAP and you are using the FC protocol.
 - Set up with working volumes and qtrees (if desired)
3. **FC environments only:** If you're using a fabric connection, verify that the switch is:
 - Set up correctly.
 - Zoned.
 - Cabled correctly according to the instructions in the *Fibre Channel and iSCSI Configuration Guide*.
 - Powered on in the correct order: switches, disk shelves, storage systems, and then the host.
4. Confirm that the host and the storage system can communicate.
5. If you currently have a version of the Solaris Host Utilities, Attack Kit, or Support Kit, remove that software.

Related information

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

Fibre Channel and iSCSI Configuration Guide -

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/fc_iscsi_config_guide.pdf

Installation overview for the Solaris Host Utilities

Installing the Host Utilities involves performing several tasks.

1. Get a copy of the compressed Host Utilities file, which contains the software package for your multipathing solution and the SAN Toolkit software package.

- Download the compressed file containing the packages from the NOW site for your processor.
 - Extract the software packages from the compressed file you downloaded.
2. Install the Host Utilities software packages.
 - a. Log in as root.
 - b. Go to the directory containing the extracted software package.
 - c. Use the `pkgadd -d` command to install the Host Utilities package for your stack.
 - d. Set the driver and system parameters. You do this using the tools provided by the Host Utilities, such as the `basic_config` command to set the driver parameters, the `creat_binding` program to set persistent bindings and driver parameters, and the `mpxio_set` script to set the system parameters.

Note: You can also set the parameters manually.
 3. Complete the configuration based on your Solaris environment.

iSCSI configuration overview

If you are using the iSCSI protocol, then you must perform some additional configuration to get it set up correctly for your environment.

1. Record the host's iSCSI node name.
2. Configure the initiator with the IP address for each storage system. You can use static, ISNS, or dynamic discovery.
3. **Veritas iSCSI environment only:** Make sure MPxIO is disabled. If you had an earlier version of the Host Utilities installed, you may need to remove the MPxIO settings that it set up and then reboot your host. To remove these settings either
 - Use the `mpxio_set -d` command to remove both the NetApp VID/PID and the the symmetric-option from the `/kernel/drv/scsi_vhci.conf` file.
 - Manually edit the `/kernel/drv/scsi_vhci.conf` file and remove the VID/PID entries.
4. (Optional) Configure CHAP on the host and the storage system.

LUN configuration overview

To complete your setup of the Host Utilities, you need to create LUNs and get the host to see them.

Configure the LUNs by performing the following tasks:

- Create at least one igroup and at least one LUN and map the LUN to the igroup.

One way to create igroups and LUNs is to use the `lun setup` command. Specify `solaris` as the value for the `ostype` attribute. You will need to supply a WWPN for each of the host's HBAs or software initiators.

- **MPxIO FC environments only:** Enable ALUA, if you haven't already done so.
- **Veritas DMP environments using LPFC drivers:** Create persistent bindings using `hbacmd`. (You can also create persistent bindings using the Host Utilities-provided `create_binding` program.)
- **Veritas DMP environments using LPFC drivers 6.11cx2 and earlier:** In the `/kernel/drv/sd.conf` file, create an entry for every NetApp LUN.
- Configure the host to discover the LUNs.
 - **LPFC drivers:** Either perform a reconfiguration reboot (`touch /reconfigure; /sbin/init 6`) 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.

You can use the `san lun` command to do this.

(Veritas DMP/FC) Information on setting up the drivers

In a Veritas DMP environment using the FC protocol, you must set up your drivers. The steps you need to perform vary based on the type drivers you are using. The following sections outline the steps for the different drivers.

Next topics

General information on getting the driver software on page 41

Emulex LPFC drivers on page 42

Sun drivers for Emulex HBAs (emlx) on page 50

Sun drivers for QLogic HBAs (qlc) on page 52

General information on getting the driver software

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

To determine which drivers are supported with the Host Utilities, check the Interoperability Matrix.

- **Emulex LPFC HBAs:** The Emulex software, including the Emulex utility programs and documentation, is available from the NetApp partner page on the Emulex site. When you are at that page, select NetApp.
- **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.
- **QLogic-branded HBAs:** The QLogic SANsurfer CLI software and documentation are available on the QLogic support site. QLogic provides a link to its NetApp partner sites. You only need this software if you have to manipulate the FCode versions on QLogic-branded HBAs for SAN booting.
- **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.

Related information

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

Emulex partner site (when this document was produced) -
<http://www.emulex.com/support/solaris/index.jsp>.

QLogic partner site (when this document was produced) -
http://support.qlogic.com/support/drivers_software.asp.

Emulex LPFC drivers

Veritas DMP supports Emulex LPFC driver versions prior to 6.21g as well as 6.21g. Some of the setup tasks you perform vary based on which driver version you have. The following sections provide instructions for working with both driver versions. You only need to look at the steps that apply to your driver version.

Next topics

[Emulex components for LPFC drivers](#) on page 42

[Downloading and extracting the Emulex software](#) on page 42

[Special Note about the Emulex LPFC driver utility `hbacmd`](#) on page 43

[Information on working with drivers prior to 6.21g](#) on page 44

[Information on working with drivers 6.21g](#) on page 47

[Determining the Emulex LPFC firmware and FCode versions](#) on page 49

[Upgrading the LPFC firmware and FCode](#) on page 49

Emulex components for LPFC drivers

The Emulex software is provided as .tar files. These files contain the components necessary for using Emulex HBAs with LPFC drivers.

The .tar files for Emulex HBAs with LPFC drivers contain the following components:

- The Host Bus Adapter driver software
- The HBA management suite, which includes the HBAnyware software and the `hbacmd` command. Sections of this document use `hbacmd` to install new firmware and to create persistent bindings. You can also perform these tasks using HBAnyware, if you prefer. For information on using HBAnyware, refer to the Emulex documentation for configuration information.
- Emulex Fibre Channel Utilities

For certain driver versions, the utility package is bundled with the driver. Beginning with the 6.21g version of the LPFC driver, Emulex started bundling a portion of the utilities with the LPFC driver itself. If you are using an LPFC driver version prior to 6.21g, the utility is a separately installed package bundle.

Downloading and extracting the Emulex software

The following steps tell you how to download and extract the Emulex software and firmware.

About this task

If your HBA uses an earlier version of the firmware than is supported by the Host Utilities, you need to download new firmware when you download the rest of the Emulex software. To determine which firmware versions are supported, check the NetApp Interoperability Matrix.

Steps

1. On the Solaris host, create a download directory for the Emulex software, such as: `mkdir /tmp/emulex`, and change to that directory.
2. To download the 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 the Emulex partner site and click the link for NetApp . Then download the required driver, firmware, FCode, utilities, and user manuals.
 - For Emulex HBAs using Sun native drivers, go to the Solaris OS download section.
3. Follow the instructions on the download page to get the driver software and place it in the `/tmp/emulex` directory you created.
4. Use the `tar xvf` 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 `emlxu_kit-<version>-sparc.tar`.

The following command lines show how to extract the software from the files.

LPFC driver bundle:

```
tar xvf solaris-HBAnyware_version-driver_version.tar
```

Emulex Utility bundle for use with Sun Native drivers:

```
tar xvf solaris-HBAnyware_version-utility_version-subversion.tar
```

Related information

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

Special Note about the Emulex LPFC driver utility hbacmd

The LPFC driver utility `hbacmd` is located in two different places.

The version you download depends on your driver version:

- For lpfc driver version 6.21g, the full path is `/opt/HBAnyware/hbacmd`.
- For lpfc driver versions prior to 6.21g, the full path is `/usr/sbin/hbanyware/hbacmd`.

This document does not use the full path in the examples. Please use the appropriate path based on the driver version you are installing.

Information on working with drivers prior to 6.21g

The following sections provide information on working with Emulex LPFC drivers prior to 6.21g.

Next topics

[Drivers prior to 6.21g: Information on installing the Emulex utilities for LPFC drivers](#) on page 44

[Drivers prior to 6.21g: Information on adding Emulex LPFC driver packages](#) on page 44

[Unbinding the HBA](#) on page 47

[Drivers prior to 6.21g: Information on removing the Emulex utilities](#) on page 45

[Drivers prior to 6.21g: Information on adding the HBAnyware package](#) on page 45

[Completing the driver installation](#) on page 47

Drivers prior to 6.21g: Information on installing the Emulex utilities for LPFC drivers

The Emulex *FCA Utilities Reference Manual* contains information on how to install the Emulex utilities for LPFC drives.

After you extract the Emulex Utilities software bundle, follow the instructions in that manual to install the package and then use the utilities.

Drivers prior to 6.21g: Information on adding Emulex LPFC driver packages

The Emulex user's manual that you downloaded contains instructions on how to install the LPFC driver package.

After you extract the driver software bundle, follow those instructions to install the LPFC driver package only.

- Do **not** install the HBAnyware packages at this time.
- Do **not** reboot the host at this time, even if the manual instructs you to do so.

Unbinding the HBA

You must unbind the Emulex HBA from the emlxs driver.

Caution: When you change the binding to the LPFC driver, the device paths for any LUNs attached to the HBA are changed, which could invalidate any `/etc/vfstab` entries that use those devices.

Steps

1. Unbind the HBA from the emlxs driver. Enter:

```
/opt/EMLXemlxu/bin/emlxdrv clear_emlxs
```

The emlxs bindings are cleared.

2. Bind the HBA to the LPFC driver. Enter:

```
/opt/EMLXemlxu/bin/emlxdrv set_lpfc_nonsun
```

The LPFC driver is bound to the HBAs.

3. Enter q to quit the emlxdrv tool.

Drivers prior to 6.21g: Information on removing the Emulex utilities

There is a potential conflict between some versions of HBAnyware and the Emulex utilities. NetApp recommends that you uninstall the Emulex utilities before you continue.

Follow the instructions in the Emulex *FCA Utilities Reference Manual* to remove the package.

Drivers prior to 6.21g: Information on adding the HBAnyware package

The Emulex user's manual that you downloaded contains instructions on how to add the HBAnyware package.

Follow those instructions to install the HBAnyware software package.

Completing the driver installation

To finish the driver installation, you must shut down the host and ensure the adapters are in set_sd_boot mode.

If the drivers are not in set_sd_boot mode, they will not bind to the LPFC driver when you boot up.

Steps

1. Set auto-boot? to false to prevent the host from booting up automatically. Enter:

```
eeeprom auto-boot?=false
```

The PROM variable auto-boot? is set to false.

2. Halt the system. Enter:

```
/sbin/init 0
```

The host is taken down to the OK> prompt.

3. Reset the host. Enter:

```
OK> reset-all
```

The host resets and runs the power-on self test. It stops at the OK> prompt.

4. Run the show-devs command to find the device paths for the HBAs.

Note: When you run this command, be sure you document the paths for your HBAs.

Running this command displays a list of the HBA paths appears. Be sure you document the paths for your HBAs. These paths could show up as `lpfc@`, `emlxs@`, or `fibre-channel@`. If they show up as `emlxs@`, continue with step 5; otherwise you can skip to step 8.

In this example, the `show-devs` command produces output similar to the following.

```
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
```

5. Switch the HBA to set-sd-boot mode. Enter:

```
OK> select <device path>
```

```
OK> set-sd-boot
```

6. Repeat step 5 for each HBA.
7. Reset the host so the boot mode changes are implemented. Enter:

```
OK> reset-all
```

The host resets and runs the power on self test. It stops at the `OK>` prompt.

8. Re-enable auto boot. Enter:

```
OK> setenv auto-boot? true
```

Auto Boot mode is turned back on.

9. Boot the host. Enter:

```
OK> boot -r
```

The host performs a reconfiguration boot.

Information on working with drivers 6.21g

The following sections provide information on working with Emulex LPFC drivers 6.21g.

Next topics

[Drivers 6.21g: Information on adding Emulex LPFC and HBAnyware driver packages](#) on page 47

[Unbinding the HBA](#) on page 47

[Completing the driver installation](#) on page 47

Drivers 6.21g: Information on adding Emulex LPFC and HBAnyware driver packages

The Emulex user's manual that you downloaded contains instructions on how to install the LPFC driver and HBAnyware package.

Use the instructions in this manual to install these features.

Do **not** reboot the host at this time, even if the Emulex manual instructs you to do so.

Unbinding the HBA

You must unbind the Emulex HBA from the emlxs driver.

Caution: When you change the binding to the LPFC driver, the device paths for any LUNs attached to the HBA are changed, which could invalidate any `/etc/vfstab` entries that use those devices.

Steps

1. Unbind the HBA from the emlxs driver. Enter:

```
/opt/EMLXemlxu/bin/emlxdrv clear_emlxs
```

The emlxs bindings are cleared.

2. Bind the HBA to the LPFC driver. Enter:

```
/opt/EMLXemlxu/bin/emlxdrv set_lpfc_nonsun
```

The LPFC driver is bound to the HBAs.

3. Enter `q` to quit the emlxdrv tool.

Completing the driver installation

To finish the driver installation, you must shut down the host and ensure the adapters are in `set_sd_boot` mode.

If the drivers are not in `set_sd_boot` mode, they will not bind to the LPFC driver when you boot up.

Steps

1. Set auto-boot? to false to prevent the host from booting up automatically. Enter:

```
eeeprom auto-boot?=false
```

The PROM variable auto-boot? is set to false.

2. Halt the system. Enter:

```
/sbin/init 0
```

The host is taken down to the OK> prompt.

3. Reset the host. Enter:

```
OK> reset-all
```

The host resets and runs the power-on self test. It stops at the OK> prompt.

4. Run the show-devs command to find the device paths for the HBAs.

Note: When you run this command, be sure you document the paths for your HBAs.

Running this command displays a list of the HBA paths appears. Be sure you document the paths for your HBAs. These paths could show up as lpfc@, emlxs@, or fibre-channel@. If they show up as emlxs@, continue with step 5; otherwise you can skip to step 8.

In this example, the show-devs command produces output similar to the following.

```
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
```

5. Switch the HBA to set-sd-boot mode. Enter:

```
OK> select <device path>
```

```
OK> set-sd-boot
```

6. Repeat step 5 for each HBA.
7. Reset the host so the boot mode changes are implemented. Enter:

```
OK> reset-all
```

The host resets and runs the power on self test. It stops at the OK> prompt.

8. Re-enable auto boot. Enter:

```
OK> setenv auto-boot? true
```

Auto Boot mode is turned back on.

9. Boot the host. Enter:

```
OK> boot -r
```

The host performs a reconfiguration boot.

Determining the Emulex LPFC firmware and FCode versions

Make sure you are using the Emulex firmware recommended for the Host Utilities when using LPFC.

About this task

To determine which version of firmware you should be using and which version you are actually using, complete the following steps.

Steps

1. Check the NetApp Interoperability Matrix to determine the current firmware requirements.
2. Use `hbacmd` to determine the version of the installed adapters. Enter:

```
hbacmd listhbas
```

The software displays the list of installed adapters and their World Wide Port Names (WWPN).

3. Use `hbacmd` to determine the firmware of each adapter. Enter:

```
hbacmd hbaattributes<adapter_wwpn>
```

The software displays information about the adapter.

Upgrading the LPFC firmware and FCode

If you are not using the Emulex firmware recommended for the Host Utilities, you must upgrade your firmware.

Before you begin

You can download the latest firmware from the NetApp page on the Emulex website.

Steps

1. Use `hbacmd` to document the current installed adapters. Enter:

```
hbacmd listhbas
```

The software displays the list of installed adapters and their WWPNs.

2. Use `hbacmd` to update the firmware. You must use this command with each adapter. Enter:

```
hbacmd download <adapter_wwpn> <firmware_filename>
```

The firmware is upgraded on the adapter.

3. Use `hbacmd` to update the FCode. This command should be repeated for each adapter. Enter:

```
hbacmd download<adapter_wwpn><fcode_filename>
```

The FCode is upgraded on the adapter.

Related information

Emulex's NetApp page is available at <http://www.emulex.com/downloads/netapp.html>

Sun drivers for Emulex HBAs (emlxs)

The Veritas DMP environment of the Host Utilities supports Emulex HBAs with Sun native drivers. The Emulex software for these drivers is provided as .tar files. You need the .tar files containing the Emulex Fibre Channel Adapter (FCA) Utilities (EMLXemlxu).

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

The sections that follow contain information on what you need to do to set up these drivers for the Host Utilities' Veritas DMP environment.

Next topics

[Installing the EMLXemlxu utilities](#) on page 50

[Determining the Emulex firmware and FCode versions for native drivers](#) on page 51

[Upgrading the firmware for native drivers](#) on page 51

[Updating your FCode HBAs with native drivers](#) on page 52

Installing the EMLXemlxu utilities

After you extract the EMLXemlxu utilities, you must install the EMLXemlxu package.

Step

1. 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 Emulex firmware and FCode versions for native drivers

Make sure you are using the Emulex firmware recommended for the Host Utilities when using LPFC.

About this task

To determine which version of firmware you should be using and which version you are actually using, complete the following steps:

Steps

1. Check the NetApp Interoperability Matrix to determine the current firmware requirements.
2. Run the `emlxadm` utility. Enter:

```
/opt/EMLXemlxu/bin/emlxadm
```

The software displays a list of available adapters.

3. Select the device that you want to check.

The software displays a menu of options.

4. Exit the `emlxadm` utility by entering `q` at the `emlxadm>` prompt.

Upgrading the firmware for native drivers

If you are not using the Emulex firmware recommended for the Host Utilities using native drivers, you must upgrade your firmware.

About this task

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

Steps

1. Run the `emlxadm` utility. Enter:

```
/opt/EMLXemlxu/bin/emlxadm
```

The software displays a list of available adapters.

2. At the `emlxadm>` prompt, enter:
`download_fw <filename>`
The firmware is loaded onto the selected adapter.
3. Exit the `emlxadm` utility by entering `q` at the `emlxadm>` prompt.
4. Reboot your host.

Updating your FCode HBAs with native drivers

If you are not using the correct FCode for HBAs using native drivers, you must upgrade it.

Steps

1. Run the `emlxadm` utility. Enter:
`/opt/EMLXemlxu/bin/emlxadm`
The software displays a list of available adapters.
2. Select the device you want to check.
The software displays a menu of options.
3. At the `emlxadm>` prompt, enter:
`download_fcode <filename>`
The FCode is loaded onto the selected adapter.
4. Exit the `emlxadm` utility by entering `q` at the `emlxadm>` prompt

Sun drivers for QLogic HBAs (qlc)

The Veritas DMP environment of the Host Utilities supports QLogic-branded and Sun-branded QLogic OEM HBAs that use the native driver (qlc) software. The following sections provide information on setting up these drivers.

Next topics

[Downloading and extracting the QLogic software](#) on page 52

[Installing the SANsurfer CLI package](#) on page 53

[Determining the FCode on QLogic cards](#) on page 53

[Upgrading the QLogic FCode](#) on page 54

Downloading and extracting the QLogic software

If you are using QLogic drivers, you must download and extract the QLogic software and firmware.

Steps

1. On the Solaris host, create a download directory for the QLogic software. Enter:
`mkdir /tmp/qlogic`
2. To download the SANsurfer CLI software, go to the QLogic website (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 multibios 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:
`uncompress scli-<version>.SPARC-X86.Solaris.pkg.Z`

Installing the SANsurfer CLI package

After you extract the QLogic files, you need to install the SANsurfer CLI package.

Steps

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 the QLogic FCode” on page 44.

Determining the FCode on QLogic cards

If you are not using the FCode recommended for the Host Utilities, you must upgrade it.

Steps

1. Check the NetApp Interoperability Matrix to determine the current FCode requirements.
2. Run the scli utility to determine whether your FCode is current or needs updating. Enter:

```
/usr/sbin/scli
```

The software displays a menu.

3. Select option 3 (HBA Information Menu).

The software displays the HBA Information Menu.

4. Select option 1 (Information).
The software displays a list of available ports.
5. Select the adapter port for which you want information.
The software displays information about that HBA port.
6. Write down the FCode version and press Return.
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.
The software displays the main menu.
8. To exit the scli utility, select option 13 (Quit).

Upgrading the QLogic FCode

If you are not using the correct FCode for HBAs using QLogic, you must upgrade it.

Steps

1. Run the scli utility. Enter:
`/usr/sbin/scli`
The software displays a menu.
2. Select option 8 (HBA Utilities).
The software displays a menu.
3. Select option 3 (Save Flash).
The software displays a list of available adapters.
4. Select the number of the adapter for which you want information.
The software displays a file name to use.
5. Enter the name of the file into which you want to save the flash contents.
The software backs up the flash contents and then waits for you to press Return.
6. Press Return.
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. .
8. When you have finished upgrading the adapters, select option 0 to return to the main menu.
9. Select option 8 (HBA Utilities).

The software displays a menu.

- 10.** Select option 1 (Update Flash).

The software displays a menu of update options.

- 11.** Select option 1 (Select an HBA Port)

The software displays a list of available adapters.

- 12.** Select the appropriate adapter number.

The software displays a list of Update ROM options.

- 13.** Select option 1 (Update Option ROM).

The software requests a file name to use.

- 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

The software upgrades the FCode.

- 15.** Press Return.

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).

The Solaris Host Utilities installation process

The Solaris Host Utilities installation process involves several tasks. You must make sure your system is ready for the Host Utilities, download the correct copy of the Host Utilities installation file, and install the software. The following sections provide information on tasks making up this process.

Next topics

[Key steps involved in setting up the Host Utilities](#) on page 57

[The software packages](#) on page 58

[Downloading the Host Utilities software from the NOW site](#) on page 58

[Installing the Solaris Host Utilities software](#) on page 59

Key steps involved in setting up the Host Utilities

Setting up the Host Utilities on your system involves both installing the software package for your stack and then performing certain configuration steps based on your stack.

Before you install the software, confirm the following:

- Your host system meets requirements and is set up correctly. Check the interoperability matrix to determine the current hardware and software requirements for the Host Utilities.
- **(Veritas DMP)** If you are using a Veritas environment, make sure Veritas is set up. In addition, to use it with the Host Utilities, you must also install the Symantec Array Support Library (ASL) and Array Policy Module (APM) for NetApp storage systems.
- You do not currently have a version of the Solaris Host Utilities, Solaris Attach Kit, or the iSCSI Support kit installed. If you previously installed one of these kits, you must remove it before installing a new kit.
- You have a copy of the Host Utilities software. You can download a compressed file containing the Host Utilities software from the NOW site.

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:

Note: (Veritas DMP/FC) If you are using a Veritas DMP environment with the FC protocol, you can use tools from the HBA Management Suite. In that case, you do not need to install the Solaris Host Utilities.

- **(MPxIO)** `mpxio_set` can be used to set or remove the symmetric option. If the host was previously set up to use the symmetric option (iSCSI) then use the `mpxio_set` tool to remove the symmetric options for the `/kernel/drv/scsi_vhci.conf` file.

- **(MPxIO and Veritas DMP)** `basic_config` configures the host system parameters. The way you use `basic_config` depends on whether you are using a Veritas DMP environment or an MPxIO (or no multipathing) environment.
- **(Veritas DMP/LPFC)** `create_binding` creates persistent bindings for systems using LPFC drivers.

After you install the Host Utilities software, you will need to configure the host system parameters. The configuration steps you perform depend on which environment you are using:

- Veritas DMP
- MPxIO

In addition, if you are using the iSCSI protocol, you must perform some additional setup steps.

The software packages

There are two Host Utilities software distribution packages.

You only need to install the file that is appropriate for your system. The two packages are:

- **SPARC processor systems:** Install this software package if you have either a Veritas DMP environment or an MPxIO environment using a SPARC processor.
- **x86/64 systems:** Install this software package if you have either a Veritas iSCSI environment or a MPxIO environment that is using an x86/64 processor.

Downloading the Host Utilities software from the NOW site

You can download the Host Utilities software package for your environment from the NOW site.

Steps

1. Log in to the NetApp NOW site and select the Download Software page.
2. Follow the prompts to reach the Software Download page.
3. Download the compressed file for your processor to a local directory on your host.

Note: If you download the file to a machine other than the Solaris host, you must move it to the host.

After you finish

Next you need to uncompress the software file and then use a command such as Solaris' `pkgadd` to add the software to your host.

Related information

NetApp NOW site - <http://now.netapp.com/>

Installing the Solaris Host Utilities software

Installing the Host Utilities involves uncompressing the files and adding the correct software package to your host.

Before you begin

Make sure you have downloaded the compressed file containing the software package for the Host Utilities.

In addition, it is a good practice to check the Solaris Host Utilities Release Notes to see if there have been any changes or new recommendations for installing and using the Host Utilities since this installation guide was produced.

Steps

1. Log in to the host system as root.
2. Place the compressed file for your processor in a directory on your host and go to that directory.

At the time this documentation was prepared, the compressed files were called:

- SPARC CPU: `netapp_solaris_host_utilities_5_1_sparc.tar.gz`
- x86/x64 CPU: `netapp_solaris_host_utilities_5_1_amd.tar.gz`

Note: The actual file names for the Host Utilities software may be slightly different from the ones shown in these steps. These are provided as examples of what the filenames look like and to use in the examples that follow. The files you download are correct.

If you are installing the `netapp_solaris_host_utilities_5_1_sparc.tar.gz` file on a SPARC system, you might put it in the `/tmp` directory on your Solaris host.

The following example places the file in that directory and then moves to the directory:

```
# cp netapp_solaris_host_utilities_5_1_sparc.tar.gz /tmp
# cd /tmp
```

3. Unzip the file using the `gunzip` command.

The software unzips the `tar.gz` files. For example, you might enter the following command like to unzip files for a SPARC system.

The following example unzips files for a SPARC system.

```
# gunzip netapp_solaris_host_utilities_5_1_sparc.tar.gz
```

4. Untar the file. You can use the `tar xvf` command to do this.

The Host Utilities scripts are extracted to the default directory.

The following example uses the `tar xvf` command to extract the Solaris installation package for a SPARC system.

```
# tar xvf netapp_solaris_host_utilities_5_1_sparc.tar
```

5. Add the packages that you extracted from tar file to your host. You can use the `pkgadd` command to do this.

The packages are added to the `/opt/NTAP/SANToolkit/bin` directory.

The following example uses the `pkgadd` command to install the Solaris installation package.

```
# pkgadd -d ./NTAPSANTool.pkg
```

6. Confirm that the toolkit was successfully installed by using the `pkginfo` command or the `ls -al` command.

This example uses the `ls -al` command to confirm that the toolkit was successfully installed.

```
[20] 11:00:50 (root@shu05) /opt/NTAP/SANToolkit/bin
# ls -la
total 3458
drwxrwxr-x  2 root    sys      512 Nov 13 11:00 ./
drwxrwxr-x  3 root    sys      512 Nov 13 11:00 ../
-r-xr-xr-x  1 root    sys    21146 Nov 10 16:28 basic_config*
-r-xr-xr-x  1 root    sys     4733 Nov 10 16:28 brocade_info*
-r-xr-xr-x  1 root    sys     5229 Nov 10 16:28 cisco_info*
-r-xr-xr-x  1 root    sys     9543 Nov 10 16:28 controller_info*
-r-xr-xr-x  1 root    sys    52670 Nov 10 16:28 create_binding*
-r-xr-xr-x  1 root    sys     9543 Nov 10 16:28 filer_info*
-r-xr-xr-x  1 root    sys    43123 Nov 10 16:28 FTP.pm*
-r-xr-xr-x  1 root    sys     5287 Nov 10 16:28 mcdata_info*
-r-xr-xr-x  1 root    sys    14052 Nov 10 16:27 mpxio_set*
-r-xr-xr-x  1 root    sys     6813 Nov 10 16:28 qllogic_info*
-r-xr-xr-x  1 root    sys      996 Nov 10 16:28 san_version*
-r-xr-xr-x  1 root    sys   1394708 Nov 10 16:28 sanlun*
-r-xr-xr-x  1 root    sys    16032 Nov 10 16:28 SHsupport.pm*
-r-xr-xr-x  1 root    sys    20169 Nov 10 16:28 solaris_info*
-r-xr-xr-x  1 root    sys    12772 Nov 10 16:28 ssd_config.pl*
-r-xr-xr-x  1 root    sys   130743 Nov 10 16:28 Telnet.pm*
-r-xr-xr-x  1 root    sys      677 Nov 10 16:28 vidpid.dat*
```

After you finish

To complete the installation, you must configure the host parameters for your environment:

- Veritas DMP

- MPxIO

If you are using iSCSI, you must also configure the initiator on the host.

Information on upgrading or removing the Solaris Host Utilities

You can easily upgrade the Solaris Host Utilities to a new version or remove an older version. If you are removing the Host Utilities, the steps you perform vary based on the version of the Host Utilities or Attach Kit that is currently installed. The following sections provide information on upgrading and removing the Host Utilities.

Next topics

[Upgrading the Solaris Host Utilities or reverting to another version](#) on page 63

[Methods for removing the Solaris Host Utilities](#) on page 63

[Uninstalling Solaris Host Utilities 5.x, 4.x, 3.x](#) on page 64

[Uninstalling the Attach Kit 2.0 software](#) on page 65

Upgrading the Solaris Host Utilities or reverting to another version

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

Steps

1. Use the Solaris `pkgrm` command to remove the Host Utilities software package you no longer need.

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 `pkgadd` command to add the appropriate Host Utilities software package.

Methods for removing the Solaris Host Utilities

You should never install the Host Utilities on a host that currently has a version of the Solaris Host Utilities or Attach Kit installed.

There are two standard methods for uninstalling the Host Utilities or Attach Kit from your system. The method you use depends on the version of the kit that is installed.

- For Solaris Host Utilities 5.x, 4.x, or 3.x, use the `pkgrm` command to remove the software package.
- For Solaris Attach Kit 2.0, use the `uninstall` script included with the Attach Kit to uninstall the software package.

Uninstalling Solaris Host Utilities 5.x, 4.x, 3.x

If you have the Solaris Host Utilities 5.x, 4.x, or 3.0 installed, you can use the `pkgrm` command to remove the software. If you want to revert to the saved parameter values, you must perform additional steps.

Steps

1. If you want to remove the parameters that were set when you ran the `basic_config` command or that you set manually after installing the Host Utilities and restore the previous values, you can do one of the following:
 - Replace the system files with the backup files you made before changing the values.
 - **(Sun native drivers)** SPARC systems and systems: `/kernel/drv/ssd.conf`.
 - **(Sun native drivers)** x86/64 systems: `/kernel/drv/sd.conf`.
 - **(Veritas DMP/LPFC drivers)** Replace `/etc/system`, `/kernel/drv/lpfc.conf`, and `/kernel/drv/sd.conf`
 - Use the `basic_config` command to revert to the saved values.

Note: You can only do this once.

 - **(Sun native drivers)** SPARC systems: execute `basic_config -ssd_unset`
 - **(Sun native drivers)** x86/64 systems: execute `basic_config -sd_unset`
 - **(Veritas DMP/LPFC drivers)** Execute `basic_config -r`
2. Use the `pkgrm` command to remove the Solaris Host Utilities software from the `/opt/NTAPSANToolkit/bin` directory.

The following command line removes the Host Utilities software package.

```
# pkgrm NTAPSANTool
```

3. To enable the changes, reboot your system.

You can use the following commands to reboot your system.

```
# touch /reconfigure
# init 6
```

4. **(MPxIO)** You can disable MPxIO by using the `stmsboot` command.

This command is only available with the following systems:

- SPARC
- x86/x64 running Solaris 10 Update 4 and newer

The following command line disables MPxIO.

```
# stmsboot -d
```

Note: The `stmsboot` command will reboot the host.

Uninstalling the Attach Kit 2.0 software

If you have the Solaris Attach Kit 2.0 installed, complete the following steps to remove the software.

Steps

1. Ensure that you are logged in as root.
2. Locate the Solaris Attach Kit 2.0 software. By default, the Solaris Attach Kit is installed in `/opt/NTAPsanlun/bin`.
3. From the `/opt/NTAPsanlun/bin` directory, enter the `./uninstall` command to remove the existing software.

You can use the following command to uninstall the existing software.

```
# ./uninstall
```

Note: The `uninstall` script automatically creates a backup copy of the `/kernel/drv/lpfc.conf` and `sd.conf` files as part of the `uninstall` procedure. It is a good practice, though, to create a separate backup copy before you begin the `uninstall`.

4. At the prompt “Are you sure you want to uninstall lpfc and sanlun packages?” enter `y`.

The `uninstall` script creates a backup copy of the `/kernel/drv/lpfc.conf` and `sd.conf` files to `/usr/tmp` and names them:

- `lpfc.conf.save`
- `sd.conf.save`

If a backup copy already exists, the `install` script prompts you to overwrite the backup copy.

5. Reboot your system.

You can use the following commands to reboot your system.

```
# touch /reconfigure  
# init 6
```

(iSCSI) Additional configuration for iSCSI environments

When you are using the iSCSI protocol, you need to perform some additional tasks to complete the installation of the Host Utilities.

You must:

- Record the host's initiator node name. You need this information to set up your storage.
- Configure the initiator with the IP address for each storage system using either static, ISNS, or dynamic discovery.
- **(MPxIO)** Run the Host Utilities' `mpxio_set` command to add the storage system's vendor ID (NetApp) and product ID (VID/PID) to the Sun MPxIO configuration file.
- **(Veritas only)** Disable MPxIO. If you have a version of the Veritas Storage Foundation that supports iSCSI, you can use it with the Host Utilities. However, because Veritas environments use DMP for multipathing, you must disable MPxIO before you can use iSCSI with Veritas.
- (Optionally) configure CHAP.

The following sections explain how to perform these tasks.

Next topics

[iSCSI node names](#) on page 67

[\(iSCSI\) Recording the initiator node name](#) on page 68

[\(iSCSI\) Storage system IP address and iSCSI static, ISNS, and dynamic discovery](#) on page 68

[\(MPxIO/iSCSI\) Configuring MPxIO with `mpxio\_set` in iSCSI environments](#) on page 69

[\(Veritas DMP/iSCSI\) Support for iSCSI in a Veritas DMP environment](#) on page 69

[\(iSCSI\) CHAP authentication](#) on page 71

iSCSI node names

To perform certain tasks, you need to know the iSCSI node name.

Each iSCSI entity on a network has a unique iSCSI node name. This is a logical name that is not linked to an IP address.

Only initiators (hosts) and targets (storage systems) are iSCSI entities. Switches, routers, and ports are TCP/IP devices only and do not have iSCSI node names.

The Solaris software initiator uses the `iqn-type` node name format:

```
iqn.yyyy-mm.backward_naming_authority:unique_device_name
```

- *yyyy* is the year and *mm* is the month in which the naming authority acquired the domain name.
- *backward_naming_authority* is the reverse domain name of the entity responsible for naming this device. An example reverse domain name is com.netapp.
- *unique_device_name* is a free-format unique name for this device assigned by the naming authority.

The following example shows a default iSCSI node name for a Solaris software initiator:

```
iqn.1986-03.com.sun:01:0003ba0da329.43d53e48
```

(iSCSI) Recording the initiator node name

You need to get and record the host's initiator node name. You use this node name when you configure the storage system.

Steps

1. On the Solaris host console, enter the following command:

```
iscsiadm list initiator-node
```

The system displays the iSCSI node name, alias, and session parameters.

2. Record the node name for use when configuring the storage system.

(iSCSI) Storage system IP address and iSCSI static, ISNS, and dynamic discovery

The iSCSI software initiator needs to be configured with one IP address for each storage system. You can use static, ISNS, or dynamic discovery.

When you enable dynamic discovery, the host uses the iSCSI `SendTargets` command to discover all of the available interfaces on a storage system. Be sure to use the IP address of an interface that is enabled for iSCSI traffic.

Note: See the Solaris Host Utilities Release Notes for issues with regard to using dynamic discovery.

Follow the instructions in the *Solaris System Administration Guide: Devices and File Systems* to configure and enable iSCSI `SendTargets` discovery. You can also refer to the `iscsiadm` man page on the Solaris host.

(MPxIO/iSCSI) Configuring MPxIO with mpzio_set in iSCSI environments

In iSCSI environments, run the Host Utilities' `mpzio_set` command to configure MPxIO. This command adds the storage system's vendor ID (NetApp) and product ID (VID/PID) to the Sun StorageTek Traffic Manager configuration file.

About this task

The format of the entries in this file is very specific. Using the `mpzio_set` command to add the required text ensures that the entry is correctly entered.

The `mpzio_set` script is located in the `/opt/NTAP/SANToolkit/bin` directory.

Steps

1. Configure MPxIO for iSCSI by entering the `mpzio_set -e` command.

The `mpzio_set -e` command adds the following four lines to the file `/kernel/drv/scsi_vhci.conf`:

```
#Added by NetApp to enable MPxIO
device-type-scsi-options-list =
"NETAPP LUN", "symmetric-option";
symmetric-option = 0x1000000;
```

2. To enable these changes, reboot your system.

Note: These changes will not take effect until you reboot your system.

You can use the following commands to reboot your system.

```
# touch /reconfigure
# init 6
```

(Veritas DMP/iSCSI) Support for iSCSI in a Veritas DMP environment

The Host Utilities support iSCSI with certain versions of Veritas.

Check the interoperability matrix to determine whether your version of Veritas supports iSCSI.

To use iSCSI with Veritas DMP, make sure that MPxIO is disabled. If you previously ran the Host Utilities on the host, you may need to remove the MPxIO settings in order to allow Veritas DMP to provide multipathing support.

Related information

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

(Veritas DMP/iSCSI) Disabling MPxIO when using iSCSI with Veritas DMP

To use iSCSI with a version of Veritas that supports the protocol, you must disable MPxIO, if it is enabled.

Before you begin

Check the interoperability matrix to confirm that your version of Veritas supports iSCSI.

About this task

You only need to perform these steps if you previously ran the Host Utilities in an MPxIO environment. In that environment, the Host Utilities configures certain settings to work with MPxIO. To use Veritas DMP with iSCSI, you must disable those settings.

Steps

1. Use one of the following methods to disable MPxIO.

- Execute the `mpxio_set` script to remove the settings. This script removes both the NetApp VID/PID and the symmetric-option from the `/kernel/drv/scsi_vhci.conf` file. Enter the command:
`mpxio_set -d`
- Edit the `/kernel/drv/scsi_vhci.conf` file and remove the VID/PID entries. The following is an example of the text you need to remove:

```
device-type-scsi-options-list =
"NETAPP LUN", "symmetric-option";
symmetric-option = 0x1000000;
```

2. Edit the `/kernel/drv/iscsi.conf` to change the line `mpxio-disable="no"`; so that it reads:
`mpxio-disable="yes"`;
3. Reboot the host.

Related information

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

(iSCSI) CHAP authentication

If you choose, you can also configure CHAP authentication. The Solaris initiator supports both unidirectional and bidirectional CHAP.

The initiator CHAP secret value you configure on the Solaris host must be the same as the inpassword value you configured on the storage system. The initiator CHAP name must be the same as the inname value you configured on the storage system.

Note: The Solaris iSCSI initiator allows a single CHAP secret value that is used for all targets. If you try to configure a second CHAP secret, it overwrites the first value you set.

Next topics

[\(iSCSI\) Configuring bidirectional CHAP](#) on page 71

[\(iSCSI\) Data ONTAP upgrades may affect CHAP configuration](#) on page 72

(iSCSI) Configuring bidirectional CHAP

Configuring bidirectional CHAP involves several steps.

About this task

For bidirectional CHAP, the target CHAP secret value you configure on the Solaris host must be the same as the outpassword value you configured on the storage system. The target CHAP username must be set to the target's iSCSI node name on the storage system. You cannot configure the target CHAP username value on the Solaris host.

Note: Make sure you use different passwords for the inpassword value and the outpassword value.

Steps

1. Set the username for the initiator.

```
iscsiadm modify initiator-node --CHAP-name sunhostname
```

2. Set the initiator password. This password must be at least 12 characters and cannot exceed 16 characters.

```
iscsiadm modify initiator-node --CHAP-secret
```

3. Tell the initiator to use CHAP authentication.

```
iscsiadm modify initiator-node -a chap
```

4. Configure bidirectional authentication for the target.

```
iscsiadm modify target-param -B enable targetIQN
```

5. Set the target username.

```
iscsiadm modify target-param --CHAP-name filerhostname targetIQN
```

6. Set the target password. Do not use the same password as the one you supplied for the initiator password. This password must be at least 12 characters and cannot exceed 16 characters.

```
iscsiadm modify target-param --CHAP-secret targetIQN
```

7. Tell the target to use CHAP authentication

```
iscsiadm modify target-param -a chap targetIQN
```

8. Configure security on the storage system.

```
iscsi security add -i initiatorIQN -s CHAP -p initpassword -n sunhostname  
-o targetpassword -m filerhostname"
```

(iSCSI) Data ONTAP upgrades may affect CHAP configuration

In some cases, if you upgrade the Data ONTAP software running on the storage system, the CHAP configuration on the storage system is not saved.

To avoid losing your CHAP settings, run the `iscsi security add` command. You should do this even if you have already configured the CHAP settings.

(Veritas DMP/FC) Tasks for completing the setup of a Veritas DMP stack using FC

To complete the Host Utilities installation when you're using a Veritas DMP stack and the FC protocol, you must configure the system parameters.

The tasks you perform vary slightly depending on your driver.

- LPFC drivers: You must modify the `/etc/system`, `/kernel/drv/lpfc.conf`, and `/kernel/drv/sd.conf` files.
- Sun native drivers: You must modify the `/kernel/drv/ssd.conf` file.

There are two ways to modify these files:

- Manually edit the files.
- Use the `basic_config` command to modify them. This command is provided as part of the 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 processor with MPxIO. For more information, see the information on configuring an MPxIO environment.

For a complete list of the host parameters that the Host Utilities recommend you change and an explanation of why those changes are recommended, see the *Host Settings Affected by the Host Utilities* document, which is on the NOW site on the SAN/IPSAN Information Library page.

Next topics

[\(Veritas DMP/FC\) Before you configure the Host Utilities for Veritas DMP](#) on page 73

[\(Veritas DMP/FC\) About the `basic\_config` command for Veritas DMP environments](#) on page 74

[\(Veritas DMP/LPFC\) Tasks involved in configuring systems using LPFC drivers](#) on page 75

[\(Veritas DMP/native\) Tasks involved in configuring systems using native drivers](#) on page 79

Related information

[SAN/IPSAN Information Library page - http://now.netapp.com/NOW/knowledge/docs/san/](http://now.netapp.com/NOW/knowledge/docs/san/)

(Veritas DMP/FC) Before you configure the Host Utilities for Veritas DMP

Before you configure the system parameters for a Veritas DMP environment, you need to perform certain tasks.

- **(LPFC drivers only)** Enter the `fcpl show cfmodes` command on the storage system to identify the storage system's `cfmode` setting. The required values for the host configuration files vary, depending on the `cfmode` setting.
- Create your own backup of the files you are modifying:
For systems using LPFC drivers, make backups of these files:
 - `/etc/system`
 - `/kernel/drv/lpfc.conf`
 - `/kernel/drv/sd.conf`

For systems using Sun native drivers, make a backup of the `/kernel/drv/ssd.conf` file.

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

Related information

Changing the Cluster `cfmode` Setting in Fibre Channel SAN Configurations - http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/SSICFMODE_1205.pdf

(Veritas DMP/FC) About the `basic_config` command for Veritas DMP environments

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 running in a Veritas DMP environment.

You execute the `basic_config` command from the command line. The options you supply with this command depend on whether your system environment uses LPFC or Sun native drivers.

The `basic_config` command sets system values to the ones recommended for use with the Host Utilities.

- For systems using LPFC drivers, the `basic_config` command can perform the following tasks:
 - Update items in the `/kernel/drv/lpfc.conf` file to the values recommended for NetApp storage systems using LPFC drivers.
 - Update time-out parameters in the `/etc/system` and `/kernel/drv/lpfc.conf` files, based on the `cfmode` that you specify.
 - Identify your Solaris OS version and set recommended values for that version in the `/etc/system` file.
 - Make backup copies of the `/etc/system`, `/kernel/drv/lpfc.conf`, and `/kernel/drv/ssd.conf` files.
 - Display the recommended values for your approval before updating the `/etc/system` and `/kernel/drv/lpfc.conf` files.

- For systems using Sun native drivers, the `basic_config` command can update items in the `/kernel/drv/ssd.conf` file to the values recommended for NetApp storage systems using Sun native drivers.

(Veritas DMP/LPFC) Tasks involved in configuring systems using LPFC drivers

If your system uses LPFC drivers, you must perform several tasks in order to complete the configuration.

- Identify the `cfmode` setting of the storage system.
- Set `/etc/system` parameters to required values.
- Set `/kernel/drv/lpfc.conf` to the required values.
- Set up persistent bindings and create an entry in the `/kernel/drv/lpfc.conf` file.

The following sections provide more information on these tasks.

Next topics

[\(Veritas DMP/LPFC\) Identifying the `cfmode`](#) on page 75

[\(Veritas DMP/LPFC\) Values for the `lpfc.conf` variables](#) on page 76

[\(Veritas DMP/LPFC\) Values for the `/etc/system` variables on systems using LPFC](#) on page 76

[\(Veritas DMP/LPFC\) The `basic\_config` options on systems using LPFC drivers](#) on page 77

[\(Veritas DMP/LPFC\) Example: Using `basic\_config` on systems with LPFC drivers](#) on page 78

(Veritas DMP/LPFC) Identifying the `cfmode`

The values you supply for certain parameters vary depending on the storage system's `cfmode`.

About this task

To identify the `cfmode`, complete the following steps.

Steps

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
```

The storage system displays its `cfmode`.

(Veritas DMP/LPFC) Values for the `lpfc.conf` variables

You can manually set the `/kernel/drv/lpfc.conf` variables so that they match the storage system `cfmode`.

The Host Utilities recommend the following 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`

The default value for the topology parameter is 0. You can reset it to either of the following:

- `topology=2` - for switch (point to point)
- `topology=4` - for arbitrated loop

For other parameters, the Host Utilities recommended values are:

- `tgt-queue-depth=256`
- `lun-queue-depth=16`
- `scan-down=0`
- `num-iocbs=1024`
- `num-bufs=1024`
- `dqfull-throttle-up-time=4`
- `dqfull-throttle-up-inc=2`

The Solaris Host Utilities provides a best-fit setting for target and/or LUN queue depths.

NetApp provides documentation for determining host queue depths and sizing. For information on NetApp FC and iSCSI storage systems, see the *iSCSI and Fibre Channel Configuration Guide*.

(Veritas DMP/LPFC) Values for the `/etc/system` variables on systems using LPFC

You can manually set the `/etc/system` variables so that they match the storage system `cfmode`.

The Host Utilities recommend the following values:

- 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 Solaris Host Utilities, you must also set the following:

```
set sd:sd_retry_count=5
```

(Veritas DMP/LPFC) The `basic_config` options on systems using LPFC drivers

On systems using LPFC drivers, the `basic_config` command has the several options you can use.

The following are the acceptable `basic_config` command lines:

```
basic_config -h
```

```
basic_config -p
```

```
basic_config -r
```

```
basic_config -v
```

```
basic_config -m <cfmode> [-l|-s]
```

The following table explains these options.

Option	Description	Notes
-h, -help	Displays the usage page	~
-l	For direct attached configurations, sets the topology parameter in the <code>/kernel/drv/lpfc.conf</code> file to loop (topology=4).	The topology option in the <code>/kernel/drv/lpfc.conf</code> 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.

Option	Description	Notes
-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 <code>telnet</code> command to do this. 2. Execute the <code>fcpl show cfmode</code> command.
-p	Prints out the existing settings.	~
-r	Reverts changes to the <code>/kernel/drv/lpfc.conf</code> and <code>/etc/system</code> files made by the <code>basic_config</code> script. This is done by copying the backup file to the current file.	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/lpfc.conf</code> and <code>/etc/system</code> files.
-s	For switched configurations, sets the topology parameter in the <code>/kernel/drv/lpfc.conf</code> file to switch point-to-point (topology=2)	The topology option in the <code>/kernel/drv/lpfc.conf</code> 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 <code>basic_config</code> script.	~

(Veritas DMP/LPFC) Example: Using `basic_config` on systems with LPFC drivers

The following example steps you through the process of using the `basic_config` command to configure system parameters on a system that is using Emulex LPFC HBAs and single-image cfmode.

Steps

1. Check the current values for the parameters. Enter the following command:

```
basic_config -p
```

The software displays the current setting for NetApp required items in the `/etc/system` and `/kernel/drv/lpfc.conf` file.

2. Use the `basic_config -m` command line to configure `/etc/system` and `/kernel/drv/lpfc.conf` file items to the recommended setting. The sample command line shown below uses a `cfmode` of `single-image`.

```
basic_config -m single_image
```

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.

This example uses the `basic_config -m single-image` command. The command displays the variables it will reset. Then it asks you to confirm this action.

```
basic_config -m single_image

Verified Emulex driver is loaded. Continuing configuration...
/etc/system file will have these changes:
|      set sd:sd_io_time=0x78|
Emulex lpfc.conf will have these changes:
|      lun-queue-depth=16;|
|      tgt-queue-depth=256;|
|      scan-down=0;|
|      num-iocbs=1024;|
|      num-bufs=1024;|
|      dqfull-throttle-up-time=4;|
|      dqfull-throttle-up-inc=2;|
|      nodev-tmo=120;|
|      linkdown-tmo=30;|

Please confirm the changes? (y/yes/n/no) y
Please reboot your host for the changes to take effect.
```

(Veritas DMP/native) Tasks involved in configuring systems using native drivers

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

There are two ways to set this values:

- Manually enter the values by editing the file.
- Automatically set the values by running the `basic_config` command.

Next topics

[\(Veritas DMP/native\) *ssd.conf* variables for systems using native drivers](#) on page 80

[\(Veritas DMP/native\) About the *basic_config* command for native drivers](#) on page 80

[\(Veritas DMP/native\) The *basic_config* options on systems using native drivers and Veritas 5.0](#) on page 81

[\(Veritas DMP/native\) Example: Using *basic_config* on systems with native drivers and DMP](#) on page 82

(Veritas DMP/native) *ssd.conf* variables for systems using native drivers

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

Note: Versions of the Host Utilities using native drivers always use single-image cfmode. If you are using native drivers and not using single-image mode, change your mode.

The required values are:

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

The Solaris Host Utilities provides a best-fit setting for target and/or LUN queue depths.

NetApp provides documentation for determining host queue depths and sizing. For information on NetApp FC and iSCSI storage systems, see the *iSCSI and Fibre Channel Configuration Guide*.

(Veritas DMP/native) About the *basic_config* command for native drivers

The *basic_config* command lets you automatically set system values to the ones recommended for use with the Host Utilities in Veritas DMP environments using Sun native drivers.

The *basic_config* command updates the settings in the `/kernel/drv/ssd.conf` file. It also sets the VID/PID information to “NETAPP LUN”.

Run the command:

```
basic_config -ssd_set
```

This command uses the following bit mask to modify the `/kernel/drv/ssd.conf` file to add the settings required for your storage system.

```
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;
```

(Veritas DMP/native) The basic_config options on systems using native drivers and Veritas 5.0

On systems using native drivers with Veritas 5.0, the `basic_config` command has the several options you can use for configuring your system.

The `basic_config` command line has the following format:

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

The following table explains these options.

Option	Description	Notes
-h, -help	Displays the usage page	~
-v	Displays version information for the <code>basic_config</code> command.	~
--vxdmp_set	Adds the required settings to <code>/kernel/drv/ssd.conf</code> for Veritas DMP with native drivers.	For systems using Veritas DMP.
-vxdmp_set_force	Performs the same tasks as the <code>-vxdmp_set</code> option, but suppresses the confirmation prompt.	For systems using Veritas DMP.
-vxdmp_unset	Removes the entries added by the <code>-vxdmp_set</code> option	For systems using Veritas DMP. Important: The <code>basic_config</code> command 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 Veritas DMP.

(Veritas DMP/native) Example: Using `basic_config` on systems with native drivers and DMP

The following example steps you through the process of using the `basic_config` command to configure system parameters on a system that is using Sun native drivers with DMP.

Step

1. Use the `basic_config -vxdmp_set` command to configure the appropriate parameters in the `/kernel/drv/ssd.conf` file to the recommended values.

The following example shows the type of output produced by the `basic_config -vxdmp_set` command when you execute it on a SPARC system.

```
# /opt/ontapNTAP/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
```

```
The original version of the /kernel/drv/ssd.conf file has been saved  
to /kernel/drv/ssd.conf.1228864383
```

```
Please reboot now
```

(MPxIO/FC) Tasks for completing the setup of a MPxIO stack

To complete the configuration when you're using a MPxIO stack, you must modify the parameters in either `/kernel/drv/ssd.conf` or `/kernel/drv/sd.conf` and set them to the recommended values.

To set the recommended values, you can either:

- Manually edit the file for your system.
- Use the `basic_config` command to automatically make the changes.

For a complete list of the host parameters that the Host Utilities recommend you change and an explanation of why those changes are recommended, see the *Host Settings Affected by the Host Utilities* document, which is on the NOW site on the SAN/IPSAN Information Library page.

Next topics

[\(MPxIO/FC\) Before configuring system parameters on a MPxIO stack using FC](#) on page 83

[\(MPxIO/FC\) Preparing for ALUA for MPxIO in FC environments](#) on page 84

[\(MPxIO/FC\) Parameter values for systems using MPxIO with FC](#) on page 85

[\(MPxIO/FC\) About the `basic\_config` command for MPxIO environments using FC](#) on page 85

[\(MPxIO/FC\) The `basic\_config` options on systems using MPxIO with FC](#) on page 86

[Example: Executing `basic\_config` on systems using MPxIO with FC](#) on page 87

Related information

[SAN/IPSAN Information Library page - `http://now.netapp.com/NOW/knowledge/docs/san/`](http://now.netapp.com/NOW/knowledge/docs/san/)

(MPxIO/FC) Before configuring system parameters on a MPxIO stack using FC

Before you configure the system parameters on a MPxIO stack using FC, you need to perform certain tasks.

- Create your own backup of the file you are modifying:
 - `/kernel/drv/ssd.conf` for systems using SPARC processors
 - `/kernel/drv/sd.conf` for systems using x86/x64 processors

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

- If MPxIO was previously installed using the Host Utilities or a Host Attach Kit prior to 3.0.1 and ALUA was not enabled, you must remove it.

Note: The Host Utilities do not support ALUA with iSCSI. While ALUA is currently supported in the iSCSI Solaris Host Utilities 3.0 and the Solaris Host Utilities using the FC protocol, it is not supported with the iSCSI protocol for the Host Utilities 5.x, the iSCSI Solaris Host Utilities 3.0.1, or Solaris 10 Update 3.

(MPxIO/FC) Preparing for ALUA for MPxIO in FC environments

You must enable ALUA if your environment supports it.

About this task

The environments that support ALUA are:

- Ones using both the FC protocol with the Host Utilities and a version of Data ONTAP that supports ALUA.
- A version of the iSCSI Host Utilities prior to 3.0.1.

Note: In iSCSI environments, ALUA is only supported for systems running the iSCSI Solaris Host Utilities 3.0. It is not supported with systems using the iSCSI protocol and running the Host Utilities 5.x, the iSCSI Solaris Host Utilities 3.0.1, or Solaris 10 Update 3.

The *Data ONTAP Block Access Management Guide for iSCSI and FC* contains details for enabling ALUA. However, you may need to do some work on the host before you can enable ALUA. If you previously installed a version of the Host Utilities or Support Kit that set the symmetric-option in the `/kernel/drv/scsi_vhci.conf` file, you must remove that option from the file. You can use the Host Utilities' `mpxio_set` command to remove the symmetric-option. You will have to reboot your host after you do this.

Complete the following steps:

Steps

1. If the symmetric-option was set because your system was previously used in an iSCSI configuration with MPxIO, execute the command:

```
mpxio_set -d
```

This script removes the symmetric-option from the `/kernel/drv/scsi_vhci.conf` file.

2. Reboot your system.

Note: This change will not take effect until you reboot your system.

You can use the following commands to reboot your system.

```
# touch /reconfigure
# init 6
```

(MPxIO/FC) Parameter values for systems using MPxIO with FC

You can manually set the parameter values for systems using MPxIO with the FC protocol by modifying `/kernel/drv/ssd.conf` (SPARC processor systems) or `/kernel/drv/sd.conf` (x86/x64 processor systems).

Both SPARC processor systems and x86/x64 processor systems using MPxIO use the same values. The required values are:

- `throttle_max=64`
- `not_ready_retries=300`
- `busy_retries=30`
- `reset_retries=30`
- `throttle_min=8`

You must also set the VIP/PID information to “NETAPP LUN”. You can use the `basic_config` command to configure this information.

(MPxIO/FC) About the `basic_config` command for MPxIO environments using FC

The `basic_config` command lets you automatically set system values to the ones recommended for use with the Host Utilities when run in MPxIO environments using the FC protocol.

This command updates the items in `/kernel/drv/ssd.conf` or `/kernel/drv/sd.conf` to the values recommended for NetApp storage systems using Sun native drivers. It also sets the VID/PID information.

The arguments you use with the `basic_config` command vary depending on whether you are using a SPARC processor system or an x86/x64 processor system.

For SPARC processor systems Run the `basic_config -ssd_set` command. This command uses a bit mask to modify the `/kernel/drv/ssd.conf` file to add the settings required for your storage system.

```
ssd-config-list="NETAPP LUN", "netapp-ssd-config";
netapp-ssd-config=1,0x9007,64,300,30,0,0,0,0,0,0,0,0,0,0,30,0,0,8,0,0;
```

For x86/x64 processor systems Run the `basic_config -sd_set` command. This command uses a bit mask to modify the `/kernel/drv/sd.conf` file to add the settings required for your storage system.

```
sd-config-list="NETAPP LUN", "netapp-sd-config";
netapp-sd-config=1,0x9c01,64,0,0,0,0,0,0,0,0,0,0,300,30,30,0,0,8,0,0;
```

(MPxIO/FC) The `basic_config` options on systems using MPxIO with FC

On systems using Sun native drivers with MPxIO and the FC protocol, the `basic_config` command has the several options you can use.

The `basic_config` command has the following format:

- For x86/64 systems
 - `basic_config [-sd_set | -sd_unset | -sd_set_force | -sd_unset_force]`
- For SPARC systems
 - `basic_config [-ssd_set | -ssd_unset | -ssd_set_force | -ssd_unset_force]`

The primary options vary slightly based on whether you are using an x86/64 processor or a SPARC processor.

The following table explains these options and specifies which options are specific to a processor type on systems using MPxIO.

Option	Description	Notes
<code>-h, -help</code>	Displays the usage page	~
<code>-sd_help</code>	Displays the <code>/kernel/drv/sd.conf</code> help page	~
<code>--sd_set</code>	Adds the required settings to the <code>/kernel/drv/sd.conf</code> file for Sun native drivers.	For x86/x64 systems using MPxIO.
<code>-sd_set_force</code>	Performs the same tasks as the <code>-sd_set</code> option, but suppresses the confirmation prompt.	For x86/x64 systems using MPxIO.

Option	Description	Notes
-sd_unset	Removes the entries added by the -sd_set option	For x86/x64 systems using MPxIO.
-sd_unset_force	Performs the same tasks as the -sd_unset option, but suppresses the confirmation prompt.	For x86/x64 systems using MPxIO.
--ssd_set	Adds the required settings to /kernel/drv/ssd.conf or the /kernel/drv/sd.conf for Sun native drivers.	For SPARC systems using MPxIO.
-ssd_set_force	Performs the same tasks as the -ssd_set option, but suppresses the confirmation prompt.	For SPARC systems using MPxIO.
-ssd_unset	Removes the entries added by the -ssd_set option	For SPARC systems using MPxIO.
-ssd_unset_force	Performs the same tasks as the -ssd_unset option, but suppresses the confirmation prompt.	For SPARC systems using MPxIO.
-v	Displays version information for the basic_config command.	~

Example: Executing basic_config on systems using MPxIO with FC

The following example steps you through the process of using the basic_config command to configure system parameters on a system that is using MPxIO with the FC protocol.

About this task

Note: If you need to remove these changes, run the `basic_config -ssd_unset` command on SPARC systems and the `basic_config -sd_unset` command on x86/x64 systems.

Steps

1. Execute the appropriate `basic_config` command for your system. The `basic_config` command has the following format:

- For SPARC systems
 - `basic_config [-ssd_set | -ssd_unset | -ssd_set_force | -ssd_unset_force]`
- For x86/64 systems
 - `basic_config [-sd_set | -sd_unset | -sd_set_force | -sd_unset_force]`

The following example shows the type of output produced by the `basic_config -ssd_set` command when you execute it on a SPARC system.

```
# /opt/ontapNTAP/SANToolkit/bin/basic_config -ssd_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
The original version of /kernel/drv/ssd.conf has been saved
to /kernel/drv/ssd.conf.1158267336

Please reboot now
```

2. To enable MPxIO on FC systems, do one of the following:
 - For SPARC systems running Solaris 10 Update 5 and later, enter:


```
# stmsboot -D fp -e
```
 - For SPARC systems running a Solaris 10 release prior to Update 5 and later, enter:


```
# stmsboot -e
```

Note: This command is also available for x86/x64 systems running Solaris 10 update 4.
 - For x86/64 systems, MPxIO is on by default. You need to manually reboot the host to enable the FC parameters.
3. To enable the changes, reboot your system.

You can use the following commands to reboot your system.

```
# touch /reconfigure
# init 6
```

(Veritas DMP) Array Support Library and Array Policy Module

The Veritas DMP stack requires that you install and configure the Symantec Array Support Library (ASL) and Array Policy Module (APM) for NetApp storage systems. In addition it is a good practice to set the Veritas restore daemon. The following sections contain information on setting up the Veritas environment.

Next topics

[\(Veritas DMP\) About the Array Support Library and the Array Policy Module](#) on page 89

[\(Veritas DMP\) Information provided by the ASL](#) on page 90

[\(Veritas DMP\) Information on upgrading the ASL and APM](#) on page 90

[\(Veritas DMP\) What an ASL array type is](#) on page 98

[\(Veritas DMP\) The storage system's FC failover mode or iSCSI configuration and the array types](#) on page 99

[\(Veritas DMP\) How sanlun displays the array type](#) on page 99

[\(Veritas DMP\) Using VxVM to display available paths](#) on page 99

[\(Veritas DMP\) Using sanlun to obtain multipathing information](#) on page 101

[\(Veritas DMP\) Veritas environments and the fast recovery feature](#) on page 102

[\(Veritas DMP\) The Veritas DMP restore daemon requirements](#) on page 102

[\(Veritas DMP\) Information on ASL error messages](#) on page 103

(Veritas DMP) About the Array Support Library and the Array Policy Module

The Array Support Library (ASL) and Array Policy Module (APM) for NetApp storage systems are provided by Symantec.

To get the ASL and APM, you must go to the Symantec Web site and download them.

Note: These are Symantec products; therefore, Symantec provides customer support if you encounter a problem using them.

To determine which versions of the ASL and APM you need for your version of the Host Utilities, check the NetApp Interoperability Matrix Tool. This information is updated frequently. Once you know which version you need, go to the Symantec Web site and download the ASL and APM.

The ASL is a NetApp-qualified library that provides information about storage array attributes configurations 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.

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

Related information

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

(Veritas DMP) Information provided by the ASL

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 the storage.

(Veritas DMP) Information on upgrading the ASL and APM

If you have the ASL and APM installed and you need to upgrade them to a newer version, you must first remove the current versions.

You can use the `pkgrm` command to uninstall the ASL and APM.

Next topics

(Veritas DMP) ASL and APM installation overview on page 91

(Veritas DMP) Determining the ASL version on page 91

(Veritas DMP) How to obtain the ASL and APM on page 92

(Veritas DMP) Installing the ASL and APM software on page 92

(Veritas DMP) Tasks to perform before you uninstall the ASL and APM on page 94

(Veritas DMP) ASL and APM installation overview

If you are using DMP with Veritas Storage Foundation 5.0 or higher, you must install the ASL and the APM. Installing the ASL and the APM involves performing several tasks.

You must:

- Verify that your configuration meets system requirements. See the NetApp Interoperability Matrix for current information about the system requirements.
- If you currently have the ASL installed, determine the its version.
- If you need to install a newer version of the ASL and APM, remove the older versions before you install the new versions.

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

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

- Obtain the new ASL and the APM.
- Install the new version of the ASL and APM.

Related information

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

(Veritas DMP) Determining the ASL version

If you currently have the ASL installed, you should check its version.

Step

1. Use the `Veritasvxddladm listversion` command to determine the AIX version.

The following is 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
libvxibm4k.so           vm-5.0-rev-1            5.0
libvxibm6k.so           vm-5.0-rev-1            5.0
libvxibm8k.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
libvxsun3k.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
libvxvp128i024.so	vm-5.0-rev-1	5.0
libvxvp12k.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

(Veritas DMP) How to obtain the ASL and APM

The ASL and APM are available from the Symantec web site. They are not included with the Host Utilities.

To get the ASL and APM, go to the Symantec Web site and download them. For Veritas 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.

Related information

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

(Veritas DMP) Installing the ASL and APM software

To install a fresh version of the ASL and APM that you downloaded from Symantec involves several steps.

Before you begin

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

Steps

1. Log in to the VxVM system as the root user.
2. If you already have your NetApp storage configured as JBOD in your VxVM configuration, remove the JBOD support for the storage by entering:

```
vxddladm rmjbod vid=NETAPP
```

3. Check the NetApp Interoperability Matrix to determine the currently supported version of the ASL and APM. Follow the link in the matrix to the ASL/APM TechNote on the Symantec Web site
4. Install the ASL and APM according to the installation instructions provided by the ASL/APM TechNote on the Symantec Web site.

Note: You should have your LUNs set up before you install the ASL and APM.

5. If your host is connected to NetApp storage, enter the `vxddladm listenclosure all` command.

By locating the NetApp Enclosure Type in the output of this command, you can verify the installation. The output shows the model name of the storage device if you are using enclosure-based naming with VxVM.

In the example that follows, the `vxddladm listenclosure all` command shows the Enclosure Types as FAS3020. To make this example easier to read, the line for the NetApp storage is shown in **bold**.

```
# vxddladm listenclosure all
```

ENCLR_NAME	ENCLR_TYPE	ENCLR_SNO	STATUS	ARRAY_TYPE
=====				
FAS30200	FAS3020	1071852	CONNECTED	
A/P-C-NETAPP				
Disk	Disk	DISKS	CONNECTED	Disk

6. If your host is not connected to storage, use the following command:

```
vxddladm listsupport all
```

The following is a sample of the output you see when you enter the `vxddladm listsupport all` command. To make this example easier to read, the line for the NetApp storage is shown in **bold**.

```
# vxddladm listsupport all
```

LIBNAME	VID
=====	
libvxCLARii0N.so	DGC
libvxcscovrts.so	CSCOVRTS
libvxemc.so	EMC
libvxengenio.so	SUN
libvxhds9980.so	HITACHI
libvxhdsalua.so	HITACHI
libvxhdsusp.so	HITACHI
libvxhpalua.so	HP, COMPAQ
libvxibm4k.so	IBM
libvxibm6k.so	IBM
libvxibm8k.so	IBM
libvxsena.so	SENA
libvxshark.so	IBM
libvxsun3k.so	SUN
libvxsun4k.so	SUN
libvxvpath.so	IBM
libvxxp1281024.so	HP
libvxxp12k.so	HP
libvxibmsvc.so	IBM
libvxnetapp.so	NETAPP

7. Enter the `vxddladm listapm all` command to verify that the APM is installed.

This command produces the following type of output. To make this example easier to read, the line for the NetApp storage is shown in **bold**.

Filename	APM Name	APM Version	Array Types	State
dmpaa	dmpaa	1	A/A	Active
dmpaaa	dmpaaa	1	A/A-A	Active
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				
dmpupalua	dmpupalua	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

After you finish

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

- If you have Data ONTAP 7.1 or higher, it is recommended that you change the cfmodes setting of your clustered systems to single-image mode and then reconfigure your host to discover the new paths to the disk.
- On the storage system, create LUNs and map them to igroups containing the WWPNs of the host HBAs.
- On the host, discover the new LUNs and configure them to be managed by VxVM.

Related information

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

(Veritas DMP) Tasks to perform before you uninstall the ASL and APM

Before you uninstall the ASL and APM, you should perform certain tasks.

- Quiesce I/O.
- Deport the disk group.

Next topics

[\(Veritas DMP\) Example of uninstalling the ASL and the APM](#) on page 95

[\(Veritas DMP\) Example of installing the ASL and the APM](#) on page 96

(Veritas DMP) Example of uninstalling the ASL and the APM

The following is an example of uninstalling the ASL and the APM when you have Veritas Storage Foundation 5.0.

If you were actually doing this uninstall, your output would vary slightly based on your system setup. Do not expect to get identical output on your system.

```
# 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 VRTSNTAPas1

The following package is currently installed:
  VRTSNTAPas1  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 "VRTSNTAPas1"

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 "VRTSNTAPas1" 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 "VRTSNTAPas1" was successful.

```

(Veritas DMP) Example of 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 installation, your output would vary slightly based on your system setup. Do not expect to get identical output on your system.

```

# 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 "VRTSNTAPasl" [y,n,?]

y

Installing Veritas NetApp Array Support Library as "VRTSNTAPasl"

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 "VRTSNTAPasl" 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.

(Veritas DMP) What an ASL array type is

The ASL reports information about the multipathing configuration to the DDL.

It reports the configuration 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 about system management, see the *Veritas Volume Manager Administrator's Guide* for details about system management.

(Veritas DMP) The storage system's FC failover mode or iSCSI configuration and the array types

In clustered storage configurations, the array type corresponds to the storage system cfmode settings or the iSCSI configuration.

If you use the standby cfmode or iSCSI configuration, 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.

(Veritas DMP) How sanlun displays the array type

The Host Utilities' sanlun utility displays information about the array type.

You use the sanlun utility to display information about paths to LUNs on the storage system.

- When the ASL is installed and the LUN is controlled by VxVM, the output of the sanlun command displays the Multipath_Policy as A/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.

(Veritas DMP) Using VxVM to display available paths

You can use VxVM to display information about available paths to a LUN.

Steps

1. Enter the following command to view all devices:

```
vxdisk list
```

The VxVM management interface displays the vxdisk device, type, disk, group and status. It also shows which disks are managed by VxVM.

The following example shows the type of output you see with the vxdisk list command.

```
# 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

This 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:

```
vxddmpadm getsubpaths dmpnodename=<device>
```

where *<device>* is the name listed under the output of the `vxddisk list` command.

The following example displays the type of output you should see with this command.

```
# vxddmpadm getsubpaths dmpnodename=FAS30201_0
```

NAME	STATE[A]	PATH-TYPE[M]	CTRL-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
-					

3. Enter the following command to obtain path information for a host HBA:

```
vxddmpadm getsubpaths ctrl=<controller_name>
```

controller_name is the controller displayed under “CTRL-NAME” in the output of the `vxddmpadm getsubpaths dmpnodename` command you entered in Step 2.

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.

The following example displays the type of output you should see.

```
# vxddmpadm getsubpaths ctrl=c2
```

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
-					

This output has been truncated to make the document easier to read.

(Veritas DMP) Using sanlun to obtain multipathing information

You can use the Host Utilities' sanlun utility to display multipathing information.

Step

1. On the host command line, enter:

```
sanlun lun show -p all
```

sanlun 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.

The partial output displayed in the following example shows eight paths to one LUN managed by VxVM.

```
# sanlun lun show -p all
filerX:/vol/vol1/data_lun60 (LUN 60) /dev/rdisk/c2t4d60s2 (FAS30201_0)
9g (9663676416) lun state: GOOD
```

path state	path type	device filename	host HBA	target port	partner port	Multipath policy
enabled	primary	/dev/rdisk/c2t4d60s2	lpfc0	1a		A/PG-C
enabled	primary	/dev/rdisk/c2t5d60s2	lpfc0	1b		A/PG-C
enabled	primary	/dev/rdisk/c3t5d60s2	lpfc1	4a		A/PG-C
enabled	primary	/dev/rdisk/c3t4d60s2	lpfc1	4b		A/PG-C
enabled	secondary	/dev/rdisk/c2t7d60s2	lpfc0		1b	A/PG-C
enabled	secondary	/dev/rdisk/c2t6d60s2	lpfc0		1a	A/PG-C

enabled	secondary	/dev/rdisk/c3t6d60s2	lpfc1	4b	A/PG-C
enabled	secondary	/dev/rdisk/c3t7d60s2	lpfc1	4a	A/PG-C

(Veritas DMP) Veritas environments and the fast recovery feature

Whether you need to enable or disable the Veritas Storage Foundation 5.0 fast recovery feature depends on your environment.

For example, if your host is using DMP for multipathing and running Veritas Storage Foundation 5.0 with the APM installed, you must have fast recovery **enabled**.

However, if your host is using MPxIO with Veritas, then you must have fast recovery **disabled**.

For details on using fast recovery with different Host Utilities Veritas environments, see the *Solaris Host Utilities 5.1 Release Notes*.

(Veritas DMP) The Veritas DMP restore daemon requirements

It is a good practice to set the Veritas restore daemon values for the restore policy and the polling interval to the Host Utilities recommended values.

These settings determine how frequently the Veritas daemon checks paths between the host and the storage system.

The Host Utilities recommended settings for these values at the time this document was produced were a restore policy of "disabled" and a polling interval of "60".

Check the Release Notes to see if these recommendations have changed.

(Veritas DMP) Setting the restore daemon interval

You can change the value of the restore daemon interval to match the recommendation for the Host Utilities.

About this task

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. The following steps take you through the process of setting the value.

Note: To see if there are new recommendations, check the Release Notes.

Steps

1. Stop the restore daemon.

```
/usr/sbin/vxdmpadm stop restore
```

2. Change the restore daemon setting to 60.

```
/usr/sbin/vxdmpadm start restore interval=60 policy=check_disabled
```

Note: This step reconfigures and restarts the restore daemon without the need for an immediate reboot.

3. 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:

```
restore_daemon_opts="interval=300 policy=check_disabled"
```

Edit the restore daemon options so that the interval is 60 seconds:

```
restore_daemon_opts="interval=60 policy=check_disabled"
```

4. Save and exit the /etc/init.d/vxvm-sysboot file.
5. Verify the changes.

```
/usr/sbin/vxdmpadm stat restored
```

(Veritas DMP) Information on ASL error messages

The ASL error messages have different levels of severity and importance.

Normally, 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.

Message severity	Definition	Action required
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.	Call Symantec Technical Support for help.
Info	Indicates that a CLAIMED status is being returned. The device functions fully, with Veritas DMP enabled, but the results seen by the user might be other than what is expected. For example, the enclosure name might change.	Call Symantec Technical Support for help.

LUN configuration and the Solaris Host Utilities

Configuring and managing LUNs involves several tasks. Whether you are executing the Host Utilities in a Veritas DMP environment or a MPxIO environment determines which tasks you need to perform. The following sections provide information on working with LUNs in both Host Utilities environments.

Next topics

[Overview of LUN configuration and management](#) on page 105

[Tasks necessary for creating and mapping LUNs](#) on page 107

[How the LUN type affects performance](#) on page 108

[Methods for creating igroups and LUNs](#) on page 108

[Best practices for creating igroups and LUNs](#) on page 109

[\(Veritas DMP\) LPFC drivers and LUNs](#) on page 109

[\(iSCSI\) Discovering LUNs](#) on page 113

[Sun native drivers and LUNs](#) on page 113

[Labeling the new LUN on a Solaris host](#) on page 115

[Methods for configuring volume management](#) on page 117

Overview of LUN configuration and management

LUN configuration and management involves a number of tasks.

The following table summarizes the tasks for all the supported Solaris environments. If a task does not apply to all environments, then it specifies which environments it does apply to. You only need to perform the tasks that apply to your environment.

Task	Discussion
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 iSCSI and FC</i>, which is available for download from the NOW site

Task	Discussion
2. (MPxIO) Enable ALUA	If your environment supports ALUA, you must have it set up to work with igroups. To see if ALUA is set up for your igroup, use the <code>igroup show -v</code> command.
3. (Veritas DMP with 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. You create the binding with either the <code>hbacmd</code> command or the <code>create_binding</code> program.
4. (Veritas DMP with LPFC drivers prior to 6.21g) Modify the <code>sd.conf</code> 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.
5. (MPxIO, Sun native drivers with Veritas DMP) Display a list of controllers	If you are using an MPxIO stack or Sun native drivers with Veritas DMP, you need to get information about the controller before you can discover the LUNs. Use the <code>cfgadm -al</code> command to display a list of controllers.

Task	Discussion
6. Discover LUNs	(Veritas LPFC drivers) The host discovers the LUNs when you perform a reconfiguration reboot (<code>touch /reconfigure; init 6</code>) on the host. You can also discover LUNs using the <code>/usr/sbin/devfsadm</code> command. If you are using drivers 6.21g or later, you can also use the <code>hbacmd</code> command to discover LUNs. (iSCSI) When you map new LUNs to the Solaris host, run the following command on the host console to discover the LUNs and create iSCSI device links: <pre>devfsadm -i iscsi</pre> (MPxIO, Sun native drivers with Veritas) 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
7. Label LUNs, if appropriate for your system	Use the Solaris format utility to label the LUNs. For optimal performance, slices or partitions of LUNs must be aligned with the WAFL volume.
8. Configure volume management software	You must configure the LUNs so they are under the control of a volume manager (SVM, ZFS, or VxVM). Use a volume manager that is supported by your Host Utilities environment.

Related information

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

Data ONTAP documentation library page -

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

Tasks necessary for creating and mapping LUNs

Before you can work with LUNs, you must set them up.

You need to:

- Create an igroup.

Note: If you have an active/active configuration, you must create a separate igroup on each system in the configuration.

- Create one or more LUNs and map the LUNs to an igroup

How the LUN type affects performance

The value you specify for the `ostype` parameter when you create a LUN can affect performance.

For optimal performance, slices or partitions of LUNs must be aligned with the WAFL volume. To achieve optimal performance, you need to provide the correct value for `ostype` for your system. There are two values for `ostype`:

- `solaris`
- `solaris_efi`

Use the `solaris` `ostype` with UFS and VxFS file systems. When you specify `solaris` as the value for `ostype` parameter, slices or partitions of LUNs are automatically aligned with the WAFL volume.

Solaris uses a newer labeling scheme, known as EFI, for LUNs that will be 2TB or larger, ZFS volumes, and SVM volumes with disk sets. For these situations, you specify `solaris_efi` as the value for the `ostype` parameter when you create the LUN. If the `solaris_efi` `ostype` is not available, you must perform special steps to align the partitions to the WAFL volume. See the *Solaris Host Utilities Release Notes* for details.

Methods for creating igroups and LUNs

There are several 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`, `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 detailed information about creating and managing LUNs, see the *Data ONTAP Block Access Management Guide for iSCSI and FCP*.

Related information

Data ONTAP Block Access Management Guide for iSCSI and FCP -
http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml/

Best practices for creating igroups and LUNs

There are several best practices you should consider when you create igroups and LUNs.

The best practices include:

- Disable scheduled auto snapshots.
- Each igroup should map to an application and should include all the initiators used by the application to access its data. (Multiple applications may be using the same initiators.)
- Do not put LUNs in the storage system's root volume. The default root volume is `/vol/vol0`.

(Veritas DMP) LPFC drivers and LUNs

There are several tasks you need to perform when setting up LUNs for a Veritas DMP environment. The following sections provide information on those tasks.

Next topics

(Veritas DMP/LPFC) Persistent bindings on page 109
(Veritas DMP/LPFC) Methods for creating persistent bindings on page 110
(Veritas DMP/LPFC) Creating WWPN bindings with hbacmd on page 110
(LPFC drivers prior to 6.21g) Modifying the sd.conf file on page 111
(LPFC drivers prior to 6.21g) Discovering LUNs on page 111
(LPFC drivers 6.21g) Discovering LUNs on page 112

(Veritas DMP/LPFC) 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).

When you're using LPFC drivers with Veritas, you need to have persistent bindings set up. They guarantee 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 Host Utilities support persistent bindings between a WWPN on the storage system and a target ID on the host.

(Veritas DMP/LPFC) Methods for creating persistent bindings

There are several tools and methods you can use to create persistent bindings.

You can use:

- **hbacmd** — This is an installation and configuration binary that loads automatically when you install the HBAnyware package. Refer to “Creating WWPN bindings with hbacmd” for information about creating persistent bindings using hbacmd.
- **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 program** — This program comes with the Host Utilities. Refer to “Creating WWPN bindings with create_binding” .
- **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.

(Veritas DMP/LPFC) Creating WWPN bindings with hbacmd

You can use hbacmd to create persistent bindings.

About this task

Make a note of the output when you use hbacmd to create persistent bindings. You will use that information when you configure the host to discover LUNs.

Caution: Previous versions of the 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 hbacmd to create persistent bindings.

Steps

1. On the Solaris host, run hbacmd to get a list of initiators. Enter:

```
hbacmd list
```

The software displays the list of installed adapters and their WWPNs. Write down the WWPNs. The next steps require you to enter that information.

2. Run hbacmd to print out the list of targets currently seen by the initiator. Enter:

```
hbacmd allnodeinfo <initiator_wwpn>
```

A list of targets appears as well as the SCSI bus and target information for that target. Record this information so that you can use it in the next step.

3. Run hbacmd to persistently bind the target to the initiator. Enter:

```
hbacmd setpersistentbinding
<initiator_wwpn><target_wwpn><scsi_bus><scsi_target><initiator_wwpn>
```

The initiator will be persistently bound to the target using the target and bus numbers provided.

4. Repeat steps 2 and 3 for each target and initiator combination that you want to persistently bind.
5. If you have not set up the `/kernel/drv/sd.conf` file, set it up as is appropriate for your system and then reboot your system.

See the section “(Drivers prior to 6.21g) Modifying the `sd.conf` file.”

6. To save the changes you have made, perform a reconfiguration reboot of the host. You can do this with the commands:

```
touch /reconfigure
/sbin/init 6
```

(LPFC drivers prior to 6.21g) Modifying the `sd.conf` file

For drivers prior to 6.21g, you must add an entry for each new target ID in the `/kernel/drv/sd.conf` file so that the host system knows which disks to probe when the system is rebooted.

Steps

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:

```
name="sd" parent="lpfc" target=0 lun=0;
```

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 `/kernel/drv/sd.conf` file for each additional storage system or cluster and for each LUN mapped to the host.

3. On the host console, reboot the host. Enter:

```
touch /reconfigure
/sbin/init 6
```

Note: If you’re performing multiple configuration steps that require rebooting, you may want to perform all those steps and then reboot.

(LPFC drivers prior to 6.21g) Discovering LUNs

There are several methods you can use to discover new LUNs and validate that they are visible on a host with LPFC driver versions prior to 6.21g.

Step

1. To discover new LUNs, use one of the methods listed below:

- Reboot the host with the reconfigure option:

```
touch /reconfigure
/sbin/init 6
```

- Enter the following commands:
`/usr/sbin/update_drv -f sd`
`/usr/sbin/devfsadm`

Note: Occasionally the `/usr/sbin/devfsadm` command does not find LUNs. If this occurs, reboot the host with the reconfigure option.

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.

(LPFC drivers 6.21g) Discovering LUNs

There are two ways to discover new LUNs and validate that they are visible on the host with LPFC driver versions 6.21g.

About this task

You use one of the following methods:

- Reboot the host with the reconfigure option:

```
touch /reconfigure
/sbin/init 6
```

- Dynamically discover LUNs by running `hbacmd rescanluns` on each target/initiator pair.

To use `hbacmd` to discover the new LUNs, complete the following steps.

Steps

1. On the Solaris host, run `hbacmd` to get a list of initiators. Enter:

```
hbacmd list
```

The software displays the list of installed adapters and their WWPNs. Write down the WWPNs. The next steps require you to enter that information.

2. Run `hbacmd` to get the list of targets attached to the given initiator. Enter:

```
hbacmd allnodelist <initiator_wwpn>
```

The software displays the list of targets currently bound to the given initiator. Record the WWPNs of each target so that you can use it in the next step.

3. Run `hbacmd` to rescan the LUNs attached to the given initiator/target pair. Enter:

```
hbacmd rescanluns <initiator_wwpn><target_wwpn>
```

The target/initiator pair you specify is rescanned. All the discovered LUNs are reported to the SCSI driver.

4. Repeat steps 2 and 3 for each target/ initiator combination.

(iSCSI) Discovering LUNs

The method you use to discover new LUNs when you're using the iSCSI protocol depends on whether you are using iSCSI with MPxIO or Veritas DMP.

Step

1. To discover new LUNs when you are using the iSCSI protocol, execute the commands that are appropriate for your environment.

- **(MPxIO)** Enter the command

```
/usr/sbin/devfsadm -c iscsi
```

- **(Veritas)** Enter the commands:

```
/usr/sbin/devfsadm -c iscsi
```

```
/usr/sbin/vxdctl enable
```

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. You can use the `format` command to label the disk.

Note: Occasionally the `/usr/sbin/devfsadm` command does not find LUNs. If this occurs, reboot the host with the `reconfigure` option (`touch /reconfigure; /sbin/init 6`)

Sun native drivers and LUNs

There are several tasks you need to perform when using Sun native drivers and working with LUNs. The following sections provide information on those tasks.

Next topics

[\(Sun native drivers\) Getting the controller number](#) on page 114

[\(Sun native drivers\) Discovering LUNs](#) on page 114

(Sun native drivers) Getting the controller number

Before you discover the LUNs, you need to determine what the controller number of the HBA is.

About this task

You must do this regardless of whether you are using Sun native drivers with MPxIO or Veritas DMP.

Step

1. Use the `cfgadm -al` command to determine what the controller number of the HBA is. If you use the `/usr/sbin/cfgadm -c configure c x` command to discover the LUNs, you need to replace `x` with the HBA controller number.

The following example uses the `cfgadm -al` command to determine the controller number of the HBA. To make the information in the example easier to read, the key lines in the output are shown in **bold**.

```
$ 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
```

(Sun native drivers) Discovering LUNs

You must both ensure that the host discovers the new LUNs and validate that the LUNs are visible on the host.

About this task

You must do this regardless of whether you are using Sun native drivers with MPxIO or Veritas DMP.

Step

1. To discover new LUNs, enter:

```
/usr/sbin/cfgadm -c configure c x
```

where `x` 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.

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 new LUN on a Solaris host

You can use the `format` utility to format and label new LUNs. This utility is a menu-driven script that is provided on the Solaris host. It works with all the environments supported by the Host Utilities.

Steps

1. On the Solaris host, enter:

```
/usr/sbin/format
```

2. At the `format>` prompt, select the disk you want to modify
3. When the utility prompts you 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.

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.)

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

AVAILABLE DISK SELECTIONS:
    0. c3t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
```

```

/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/sd@0,0
1. c4t0d0 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128>
/pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,0
2. c4t0d1 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128>
/pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,1
3. c4t0d2 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128>
/pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,2
4. c4t0d3 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128>
/pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,3

```

Specify disk (enter its number): 1

selecting c4t0d0

[disk formatted]

...

Disk not labeled. Label it now? y

Example 2: This example labels disk number 2 on a system that uses Sun native drivers with Veritas DMP. (Portions of this example have been removed to make it easier to review.)

```
$ 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: **TESTER**

[disk formatted]

...

Disk not labeled. Label it now? y

Example 3: This example runs the `fdisk` command and then labels disk number 15 on an x86/x64 system. You must run the `fdisk` command before you can label a LUN.

```
$Specify disk (enter its number): 15
```

```
selecting C4t60A9800043346859444A2D367047492Fd0
```

```
[disk formatted]
```

```
FORMAT MENU:
```

```

disk      - select a disk
type      - select (define) a disk type
partition - select (define) a partition table
current   - describe the current disk
format    - format and analyze the disk
fdisk     - run the fdisk program
repair    - repair a defective sector
label     - write label to the disk
analyze   - surface analysis
defect    - defect list management
backup    - search for backup labels
verify    - read and display labels
save      - save new disk/partition definitions
inquiry   - show vendor, product and revision
volname   - set 8-character volume name
!>cmd>    - execute >cmd>, then return
quit
format> label
Please run fdisk first.
format> fdisk
No fdisk table exists. The default partition for the disk is:

    a 100% "SOLARIS System" partition

Type "y" to accept the default partition, otherwise type "n" to edit
the
partition table.
y
format> label
Ready to label disk, continue? y
format>

```

Methods for 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.

The tools you use to manage your volumes depend on the environment you are working in: Veritas DMP or MPxIO.

Veritas DMP If you are in a Veritas DMP environment, even if you are using Sun native drivers or the iSCSI protocol, you must use VxVM to manage the LUNs. You can use the following Veritas commands to work with LUNs:

- The Veritas `/usr/sbin/vxdctl enable` command brings new LUNs under Veritas control.
- The Veritas `/usr/sbin/vxdiskadm` utility manages existing disk groups.

MPxIO If you are in a MPxIO environment, you can manage LUNs using SVM, ZFS, or, in some cases, VxVM

Note: To use VxVM in an MPxIO environment, first check the NetApp Interoperability Matrix Tool to see if your environment supports VxVM.

For additional information, refer to the documentation that shipped with your volume management software

Related information

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

The sanlun utility

The `sanlun` utility is a tool provided by the Host Utilities that helps collect and report information about paths to your devices and how they map to LUNs on the storage system. You can also use `sanlun` to display information about the host HBAs. The following sections provide information on `sanlun`.

Next topics

[Displaying host LUN information with `sanlun`](#) on page 119

[Displaying path information with `sanlun`](#) on page 121

[Explanation of the `sanlun lun show -p` output](#) on page 122

[Displaying host HBA information with `sanlun`](#) on page 123

Displaying host LUN information with `sanlun`

You can use `sanlun` to display information about the LUNs connected to the host.

Steps

1. Ensure that you are logged in as root on the host.
2. Change to the `/opt/NTAP/SANToolkit/bin`:
`cd /opt/NTAP/SANToolkit/bin`
3. Enter the `sanlun lun show` command to display LUN information. The command has the following format:

```
sanlun lun show [-v] [-d host device filename | all | storagesystem name |  
storagesystem name:storagesystem pathname] -p
```

`-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/rdisk`.
- *storagesystem name* lists all storage system LUNs under `/dev/rdisk` on that storage system.
- *storagesystem name:storagesystem pathname* lists all storage system LUNs under `/dev/rdisk` that are connected to the storage system path name LUN on that storage system.

`-p` provides information about the primary and secondary paths available to the LUN when you are using multipathing. You cannot use the `-d` option if you use `-p`. Use the following format:

```
sanlun lun show -p [storagesystem name:storagesystem pathname | storagesystem name | all ]
```

If you enter `sanlun lun show`, `sanlun lun show -p`, or `sanlun lun show -v` without any parameters, the utility responds as if you had included the `all` parameter.

The requested LUN information is displayed.

For example, you might enter:

- `sanlun lun show -p`
to display a listing of all the paths associated with the LUN. This information is useful if you need to set up path ordering or troubleshoot a problem with path ordering.
- `sanlun lun show -d /dev/rdisk/<x>`
to display the summary listing of the LUN(s) associated with the host device `/dev/rdisk<x>` where `x` is a device such as `/dev/rdisk/c#t#d#`.
- `sanlun lun show -v all`
to display verbose output for all the LUN(s) currently available on the host.
- `sanlun lun show toaster`
to display a summary listing of all the LUNs available to the host served by the storage system called `toaster`.
- `sanlun lun show toaster:/vol/vol0/lun0`
to display a summary listing of all the LUNs available to the host served by `lun0` on `toaster`.

Note: When you specify either the `sanlun lun show <storage_system_name>` or the `sanlun lun show <storage_system_name:storage_system_pathname>` command, the utility displays only the LUNs that have been discovered by the host. LUNs that have not been discovered by the host are not displayed.

The following is an example of the output you see when you use the `sanlun show` command in verbose mode on a system using the Veritas DMP stack.

```
# ./sanlun lun show -v
filer: lun-pathname          device filename      adapter lun size
lun state
filerX: /vol/vol1/hostA_lun2  /dev/rdisk/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
```

Displaying path information with sanlun

You can use `sanlun` to display information about the paths to the storage system.

Steps

1. Ensure that you are logged in as root on the host.
2. Use the `cd` command to change to the `/opt/NTAP/SANToolkit/bin` directory:
3. At the host command line, enter the following command to display LUN information:

```
sanlun lun show -p all -v
```

`-p` provides information about the optimized (primary) and non-optimized (secondary) paths available to the LUN when you are using multipathing.

Note: (MPxIO stack) MPxIO makes the underlying paths transparent to the user. It only exposes a consolidated device such as `/dev/rdsd/c7t60A980004334686568343655496C7931d0s2`. This is the name generated using the LUN's serial number in the IEEE registered extended format, type 6. The Solaris host receives this information from the SCSI Inquiry response. As a result, `sanlun` cannot display the underlying multiple paths. Instead it displays the target port group information. You can use the `mpathadm` or `luxadm` command to display the information if you need it. .

`all` lists all storage system LUNs under `/dev/rdsd`.

`-v` turns on verbose mode. Because MPxIO provides multiple paths, `sanlun` selects one of the available paths to display. This is not a fixed path and could vary each time you execute the `sanlun` command with the `-v` option.

The following example uses the `-p` option to display information about the paths.

```
# sanlun lun show -p
ONTAP_PATH: fas3020-rtphome21:/vol/vol0/lun0
LUN: 0
LUN Size:    20g (21474836480)
Host Device: /dev/rdsd/c7t60A980004334686568343655496C7931d0s2
LUN State: GOOD  Filer_CF_State: Cluster Enabled
Multipath_Policy: Native  Multipath-provider: Sun Microsystems
TPGS flag: 0x10Filer Status: TARGET PORT GROUP SUPPORT ENABLED
Target Port Group : 0x1001
Target Port Group State: Active/optimized
Vendor unique Identifier : 0x10 (2GB FC)
```

```

Target Port Count: 0x2
    Target Port ID : 0x1
    Target Port ID : 0x2

Target Port Group : 0x3002
Target Port Group State: Active/non-optimized
Vendor unique Identifier : 0x30 (2GB FC)
Target Port Count: 0x2
    Target Port ID : 0x101
    Target Port ID : 0x102
Target Port Group : 0x3002
Target Port Group State: Active/non-optimized
Vendor unique Identifier : 0x30 (2GB FC)
Target Port Count: 0x2
    Target Port ID : 0x101
    Target Port ID : 0x102

ONTAP_PATH: fas3020-rtphome21:/vol/vol1/lun1
LUN: 1
LUN Size:      20g (21474836480)
Host Device: /dev/rdisk/c7t60A980004334686568343655496D5464d0s2
LUN State: GOOD  Filer_CF_State: Cluster Enabled
Multipath_Policy: Native  Multipath-provider: Sun Microsystems
TPGS flag: 0x10Filer Status: TARGET PORT GROUP SUPPORT ENABLED
Target Port Group : 0x1001
Target Port Group State: Active/optimized
Vendor unique Identifier : 0x10 (2GB FC)
Target Port Count: 0x2
    Target Port ID : 0x1
    Target Port ID : 0x2

Target Port Group : 0x3002
Target Port Group State: Active/non-optimized
Vendor unique Identifier : 0x30 (2GB FC)
Target Port Count: 0x2
    Target Port ID : 0x101
    Target Port ID : 0x102

```

Explanation of the sanlun lun show -p output

The `san lun lun show -p` command provides details for both MPxIO stacks and Veritas DMP stacks.

- **(MPxIO stack)** The target port group and its state, count, and ID. See information on ALUA in the *Data ONTAP® Block Access Management Guide for FCP*.
- **(MPxIO stack)** The vendor unique identifier.
- **(Veritas DMP stack)** path state—Whether the path is enabled or disabled.
- (Veritas DMP stack) path type:

- 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 when 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 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 path
- Multipathing policy—The multipathing policy is one of the following:
 - **(Veritas DMP stack) A/A (Active/Active)**—Veritas 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.
 - **(Veritas DMP stack) 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.
 - **(MPxIO stack) Native**—this is displayed as the multipathing policy when MPxIO with ALUA is used for Multipathing.

Displaying host HBA information with sanlun

You can use `sanlun` to display information about the host HBA.

Steps

1. Ensure that you are logged in as root on the host.
2. Change to the `/opt/NTAP/SANToolkit/bin` directory.
3. At the host command line, enter the following command to display host HBA information:

```
./sanlun fcp show adapter [-c] [-v] [adapter name | all]
```

`-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.

The FC adapter information is displayed.

The following command line displays information about the adapter on a system using the qlc driver.

```
# ./sanlun fcp show adapter -v

adapter name: qlc0
WWPN: 21000003ba16dec7
WWNN: 20000003ba16dec7
driver name: 20060630-2.16
model: 2200
model description: 2200
serial number: Not Available
hardware version: Not Available
driver version: 20060630-2.16
firmware version: 2.1.144
Number of ports: 1
port type: Private Loop
port state: Operational
supported speed: 1 GBit/sec
negotiated speed: 1 GBit/sec
OS device name: /dev/cfg/c1
adapter name: qlc1
WWPN: 210000e08b88b838
WWNN: 200000e08b88b838
driver name: 20060630-2.16
model: QLA2462
model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel
serial number: Not Available
hardware version: Not Available
driver version: 20060630-2.16
firmware version: 4.0.23
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/c2
adapter name: qlc2
WWPN: 210100e08ba8b838
WWNN: 200100e08ba8b838
driver name: 20060630-2.16
model: QLA2462
model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel
serial number: Not Available
hardware version: Not Available
driver version: 20060630-2.16
```

```
firmware version: 4.0.23
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/c3
```


Troubleshooting

If you encounter a problem while running the Host Utilities, here are some tips and troubleshooting suggestions.

The sections that follow contain information on:

- Best practices, such as checking the *Release Notes* to see if any information has changed.
- Suggestions for checking your system.
- Information on possible problems and how to handle them.
- Diagnostic tools that you can use to gather information about your system.

Next topics

[Check the Release Notes](#) on page 127

[About the troubleshooting sections that follow](#) on page 128

[Possible iSCSI issues](#) on page 139

[Diagnostic utilities for the Host Utilities](#) on page 141

[Possible MPxIO issues](#) on page 145

Check the Release Notes

The *Release Notes* contain the most up-to-date information on known problems and limitations. The *Release Notes* also contain information on how to look up information about known bugs.

The *Release Notes* are updated when there is new information about the Host Utilities. It is a good practice to check the *Release Notes* before you install the Host Utilities and any time you encounter a problem with the Host Utilities.

You can access the *Release Notes* from the NOW site. A link to the *Release Notes* is online in the SAN/IPSAN library. A link is also available on the NOW Description Page for the Solaris Host Utilities. You get to the Solaris Host Utilities Description Page by following the FC Host Utilities links for Solaris that are on the Download Software page.

Related information

[NOW Download Software](http://now.netapp.com/NOW/cgi-bin/software) - <http://now.netapp.com/NOW/cgi-bin/software>

[SAN/IPSAN library](http://now.netapp.com/NOW/knowledge/docs/san/) - <http://now.netapp.com/NOW/knowledge/docs/san/>

About the troubleshooting sections that follow

The troubleshooting sections that follow help you verify your system setup.

If you have any problems with the Host Utilities, make sure your system setup is correct.. As you go through the following sections, keep in mind:

- For more information on 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 *Data ONTAP Block Access Management Guide for iSCSI and FC* for your version of Data ONTAP.
- 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 the `basic_config` command and the `san lun` command, both of which are useful when checking your system.
- To make sure you have the current version of the system components, see the Interoperability Matrix. 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.

Next topics

Check the version of your host operating system on page 129

Confirm the HBA is supported on page 129

(MPxIO, native drivers) Ensure that MPxIO is configured correctly for ALUA on FC systems on page 131

(MPxIO, FC) Ensure that MPxIO is enabled on SPARC systems on page 131

(MPxIO) Ensure that MPxIO is enabled on iSCSI systems on page 132

(MPxIO) Verify that MPxIO multipathing is working on page 133

(Veritas DMP) Check that the ASL and APM have been installed on page 134

(Veritas) Check VxVM on page 134

(MPxIO) Check the Solaris Volume Manager on page 136

(MPxIO) Check settings in `ssd.conf` or `sd.conf` on page 136

Check the storage system setup on page 137

(MPxIO/FC) Check the ALUA settings on the storage system on page 137

Verify that the switch is installed and configured on page 138

Determine whether to use switch zoning on page 138

Power up equipment in the correct order on page 139

Verify that the host and storage system can communicate on page 139

Related information

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

Check the version of your host operating system

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

You can use the `cat /etc/release` command to display information about your operating system.

The following example checks the operating system version on a SPARC system.

```
# cat /etc/release
                Solaris 10 5/08 s10s_u5wos_10 SPARC
        Copyright 2008 Sun Microsystems, Inc. All Rights Reserved.
                Use is subject to license terms.
                Assembled 24 March 2008
```

The following example checks the operating system version on an x86 system.

```
# cat /etc/release
                Solaris 10 5/08 s10x_u5wos_10 X86
        Copyright 2008 Sun Microsystems, Inc. All Rights Reserved.
                Use is subject to license terms.
                Assembled 24 March 2008
```

Confirm the HBA is supported

You can use the `san lun` command to display information on the HBA and the NetApp Interoperability Matrix to determine if the HBA is supported. Supported HBAs should be installed before you install the Host Utilities.

The `san lun` command is part of the Host Utilities SAN Toolkit.

If you are using MPxIO, you can also use the `fcinfo hba-port` command to get information about the HBA.

1. The following example uses the `san lun` command to check a QLogic HBA in an environment using a Sun native qlc driver.

```
sanlun fcp show adapter -v

adapter name:      qlc1
WWPN:              210000e08b88b838
WWNN:              200000e08b88b838
driver name:       20060630-2.16
model:             QLA2462
model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel
serial number:     Not Available
hardware version:  Not Available
driver version:    20060630-2.16
firmware version:  4.0.23
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/c2

adapter name:   qlc2
WWPN:          210100e08ba8b838
WWNN:          200100e08ba8b838
driver name:    20060630-2.16
model:         QLA2462
model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel
serial number:  Not Available
hardware version: Not Available
driver version: 20060630-2.16
firmware version: 4.0.23
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/c3

```

2. The following example uses the `sanlun` command to check an HBA in an environment using LPFC drivers.

```

# sanlun fcp show adapter -v

adapter name:      lpfc0
WWPN:              10000000c9512706
WWNN:              20000000c9512706
driver name:       lpfc
model:             LP11002-M4
model description: Emulex LP11002-M4 4Gb 2port FC: PCI-X2 SFF HBA
serial number:     VM55099793
hardware version:  1036406d
driver version:    6.21g; HBAAPI(I) v2.0.f, 01-23-06
firmware version:  2.70A5 (B2F2.70A5)
Number of ports:   1
port type:         Fabric
port state:        Operational
supported speed:   4 GBit/sec
negotiated speed:  2 GBit/sec
OS device name:    /devices/pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1

adapter name:      lpfc1
WWPN:              10000000c9512707
WWNN:              20000000c9512707
driver name:       lpfc
model:             LP11002-M4
model description: Emulex LP11002-M4 4Gb 2port FC: PCI-X2 SFF HBA
serial number:     VM55099793
hardware version:  1036406d
driver version:    6.21g; HBAAPI(I) v2.0.f, 01-23-06
firmware version:  2.70A5 (B2F2.70A5)

```

```

Number of ports: 1
port type:      Fabric
port state:     Operational
supported speed: 4 GBit/sec
negotiated speed: 2 GBit/sec
OS device name: /devices/pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1,1

```

Related information

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

(MPxIO, native drivers) Ensure that MPxIO is configured correctly for ALUA on FC systems

Configurations using native drivers with MPxIO and FC in a NetApp clustered environment require that ALUA be enabled.

In some cases, ALUA may have been disabled on your system and you will need to reenable it. For example, if your system was used for iSCSI or was part of a single NetApp storage controller configuration, the `symmetric-option` may have been set. This option disables ALUA on the host.

To enable ALUA, you must remove the `symmetric-option` by either:

- Running the `mpxio_set -d` command. This command automatically comments out the the `symmetric-option` section. Or
- Editing the appropriate section in the `/kernel/drv/scsi_vhci.conf` file to manually comment it out. The example below displays the section you must comment out.

Once you comment out the option, you must reboot your system for the change to take effect.

The following example is a the section of the `/kernel/drv/scsi_vhci.conf` file that you must comment out if you want to enable MPxIO to work with ALUA. This section has been commented out.

```

#device-type-scsi-options-list =
#"NETAPP LUN", "symmetric-option";
#symmetric-option = 0x1000000;

```

(MPxIO, FC) Ensure that MPxIO is enabled on SPARC systems

When you use a MPxIO stack on a SPARC system, you must manually enable MPxIO. If you encounter a problem, make sure that MPxIO is enabled.

Note: On x86/64 systems, MPxIO is enabled by default.

To enable MPxIO on a SPARC system, use the `stmsboot` command. This command modifies the `fp.conf` file to set the `mpxio_disable=` option to `no` and updates `/etc/vfstab`.

After you use this command, you must reboot your system.

The options you use with this command vary depending on your version of Solaris. For systems:

- Running Solaris 10 update 5, execute: `stmsboot -D fp -e`
- Prior to Solaris 10 update 5, execute: `stmsboot -e`

For example, if MPxIO is not enabled on a system running Solaris 10 update 5, you would enter the following commands. The first command enables MPxIO by changing the `fp.conf` file to read `mpxio_disable=no`. It also updates `/etc/vfstab`. The following commands reboot the system. You must reboot the system for the change to take effect.

```
# stmsboot -D fp -e
# touch /reconfigure
# init 6
```

(MPxIO) Ensure that MPxIO is enabled on iSCSI systems

While MPxIO should be enabled by default on iSCSI systems, you can confirm this by viewing the iSCSI configuration file `/kernel/drv/iscsi.conf` file.

When MPxIO is enabled, this file has the `mpxio-disable` set to "no".

`mpxio-disable="no"`

If this line is set to "yes", you must change it by either

- Running the `mpxio_set -e` command. This command sets the symmetric-option. Or
- Editing the appropriate section in the `/kernel/drv/iscsi.conf` file to manually set the command to "no". The example below displays the section you must comment out.

You **must** reboot your system for the change to take effect.

Here is an example of a `/kernel/drv/iscsi.conf` file that has MPxIO enabled. The line that enables MPxIO, `mpxio-disable="no"` is in bold to make it easy to locate.

```
# Copyright 2006 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
#ident "@(#)iscsi.conf 1.2 06/06/12 SMI"
name="iscsi" parent="/" instance=0;
ddi-forceattach=1;
Chapter 3: Configuring the initiator 23
#
# I/O multipathing feature (MPxIO) can be enabled or disabled using
# mpxio-disable property. Setting mpxio-disable="no" will activate
# I/O multipathing; setting mpxio-disable="yes" disables the
feature.
```

```
#
# Global mpzio-disable property:
#
# To globally enable MPxIO on all iscsi ports set:
# mpzio-disable="no";
#
# To globally disable MPxIO on all iscsi ports set:
# mpzio-disable="yes";
#
mpzio-disable="no";
tcp-nodelay=1;
...
```

(MPxIO) Verify that MPxIO multipathing is working

You can confirm that multipathing is working in an MPxIO environment by using either a Host Utilities tool such as the `san lun lun show` command or a Solaris tool such as the `mpathadm` command.

The `san lun lun show` command displays the disk name. If MPxIO is working, you should see a long name similar to the following:

```
/dev/rdisk/c5t60A9800004334686568343771474A4D42d0s2
```

The long, consolidated Solaris device name is generated using the LUN's serial number in the IEEE registered extended format, type 6. The Solaris host receives this information in the SCSI Inquiry response.

Another way to confirm that MPxIO is working is to check for multiple paths. To view the path information, you need to use the `mpathadm` command. The `san lun` cannot display the underlying multiple paths because MPxIO makes these paths transparent to the user when it displays the consolidated device name shown above.

In this example, the `mpathadm list lu` command displays a list of all the LUNs.

```
# mpathadm list lu
  /dev/rdisk/c3t60A9800004334612F466F4C6B72483362d0s2
    Total Path Count: 8
    Operational Path Count: 8
  /dev/rdisk/c3t60A9800004334612F466F4C6B72483230d0s2
    Total Path Count: 8
    Operational Path Count: 8
  /dev/rdisk/c3t60A9800004334612F466F4C6B7248304Dd0s2
    Total Path Count: 8
    Operational Path Count: 8
  /dev/rdisk/c3t60A9800004334612F466F4C6B7247796Cd0s2
    Total Path Count: 8
    Operational Path Count: 8
  /dev/rdisk/c3t60A9800004334612F466F4C6B72477838d0s2
    Total Path Count: 8
    Operational Path Count: 8
```

```
/dev/rdsd/c3t60A980004334612F466F4C6B72477657d0s2
Total Path Count: 8
Operational Path Count: 8
```

(Veritas DMP) Check that the ASL and APM have been installed

You must have the ASL and APM installed in order for the Veritas DMP to identify whether the path is primary or secondary

Without the ASL and APM, the DSM treats all paths as equal, even if they are secondary paths. As a result, you may see I/O errors on the host. On the storage system, you may see the Data ONTAP error:

FCP Partner Path Misconfigured

If you encounter the I/O errors on the host or the Data ONTAP error, make sure you have the ASL and APM installed.

(Veritas) Check VxVM

The Host Utilities supports the VxVM for the Veritas DMP stack and certain configurations of the MPxIO stack. You can use the `vxdisk list` command to quickly check the VxVM disks and the `vxprint` command to view volume information.

Note: (MPxIO) To determine which versions of VxVM are supported with MPxIO, see the NetApp Interoperability Matrix.

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

This example uses the `vxdisk list` command. This output in 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
...
#
```

This next example uses the `vxprint` command to display information on the volumes. The output in this example has been truncated to make it easier to read.

```
# vxprint -g scdg
TY NAME      ASSOC      KSTATE  LENGTH  PLOFFS  STATE  TUTILO
```

PUTILO						
dg_scdg	scdg	-	-	-	-	-
-						
dm_disk00	Disk_49	-	10475520	-	-	-
-						
dm_disk01	Disk_33	-	10475520	-	-	-
-						
dm_disk02	Disk_51	-	10475520	-	-	-
-						
dm_disk03	Disk_54	-	10475520	-	-	-
-						
dm_disk04	Disk_17	-	10475520	-	-	-
-						
dm_disk05	Disk_39	-	10475520	-	-	-
-						
dm_disk06	Disk_58	-	10475520	-	-	-
-						
dm_disk07	Disk_24	-	10475520	-	-	-
-						
dm_disk08	Disk_11	-	10475520	-	-	-
-						
dm_disk09	Disk_34	-	10475520	-	-	-
-						
dm_disk10	Disk_35	-	10475520	-	-	-
-						
dm_disk11	Disk_29	-	10475520	-	-	-
-						
dm_disk12	Disk_53	-	10475520	-	-	-
-						
dm_disk13	Disk_2	-	10475520	-	-	-
-						
dm_disk14	Disk_12	-	10475520	-	-	-
-						
dm_disk15	Disk_5	-	10475520	-	-	-
-						
dm_disk16	Disk_44	-	10475520	-	-	-
-						
dm_disk17	Disk_32	-	10475520	-	-	-
-						
dm_disk18	Disk_10	-	10475520	-	-	-
-						
dm_disk19	Disk_27	-	10475520	-	-	-
-						
v_raw_vol	fsngen	ENABLED	209510400	-	ACTIVE	-
-						
pl_raw_vol-01	raw_vol	ENABLED	209510400	-	ACTIVE	-
-						
sd_disk00-01	raw_vol-01	ENABLED	10475520	0	-	-
-						
sd_disk05-01	raw_vol-01	ENABLED	10475520	10475520	-	-
-						
sd_disk10-01	raw_vol-01	ENABLED	10475520	20951040	-	-
-						
sd_disk15-01	raw_vol-01	ENABLED	10475520	31426560	-	-
-						

sd disk01-01	raw_vol-01	ENABLED	10475520	0	-	-
-	-	-	-	-	-	-
sd disk06-01	raw_vol-01	ENABLED	10475520	10475520	-	-
-	-	-	-	-	-	-
sd disk11-01	raw_vol-01	ENABLED	10475520	20951040	-	-
-	-	-	-	-	-	-
sd disk16-01	raw_vol-01	ENABLED	10475520	31426560	-	-
-	-	-	-	-	-	-
sd disk02-01	raw_vol-01	ENABLED	10475520	0	-	-
-	-	-	-	-	-	-
sd disk07-01	raw_vol-01	ENABLED	10475520	10475520	-	-
-	-	-	-	-	-	-
sd disk12-01	raw_vol-01	ENABLED	10475520	20951040	-	-
-	-	-	-	-	-	-
sd disk17-01	raw_vol-01	ENABLED	10475520	31426560	-	-
-	-	-	-	-	-	-
sd disk03-01	raw_vol-01	ENABLED	10475520	0	-	-
-	-	-	-	-	-	-
sd disk08-01	raw_vol-01	ENABLED	10475520	10475520	-	-
-	-	-	-	-	-	-
sd disk13-01	raw_vol-01	ENABLED	10475520	20951040	-	-
-	-	-	-	-	-	-
sd disk18-01	raw_vol-01	ENABLED	10475520	31426560	-	-
-	-	-	-	-	-	-

Related information

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

<http://www.ibm.com/storage/support/nas/>

(MPxIO) Check the Solaris Volume Manager

If you are using the MPxIO version of the Host Utilities with the Solaris Volume Manager (SVM), it is a good practice to check the condition of the volumes.

The `metastat -a` command lets you quickly check the condition of SVM volumes.

Note: In a Sun Cluster environment, metaset and their volumes are only displayed on the node that is controlling the storage.

See your Solaris documentation for more information on working with the SVM.

The following sample command line checks the condition of SVM volumes:

```
# metastat -a
```

(MPxIO) Check settings in `ssd.conf` or `sd.conf`

Verify that you have the correct settings in the configuration file for your system.

The file you need to modify depends on the processor your system uses:

- SPARC systems with MPxIO enabled use the `ssd.conf` file. You can use the `basic_config -ssd_set` command to update the `/kernel/drv/ssd.conf` file.
- x86/64 systems with MPxIO enabled use the `sd.conf` file. You can use the `basic_config -sd_set` command to update the `sd.conf` file.

Example of `ssd.conf` file (MPxIO on a SPARC system): You can confirm the `ssd.conf` file was correctly set up by checking it to ensure that it contains the following:

```
ssd-config-list="NETAPP LUN", "netapp-ssd-config";
netapp-ssd-config=1,0x9007,64,300,30,0,0,0,0,0,0,0,0,0,30,0,0,8,0,0;
```

Example of `sd.conf` file (MPxIO on an x86/64 system): You can confirm the `sd.conf` file was correctly set up by checking it to ensure that it contains the following:

```
sd-config-list="NETAPP LUN", "netapp-sd-config";
netapp-sd-config=1,0x9c01,64,0,0,0,0,0,0,0,0,0,0,300,30,30,0,0,8,0,0;
```

Check the storage system setup

Make sure your storage system set up correctly.

For information on what needs to be done, see the list of storage system requirements in the section “Prerequisites” on page 20

(MPxIO/FC) Check the ALUA settings on the storage system

In MPxIO environments using FC, you must have ALUA set on the storage system to work with igroups.

You can verify that you have ALUA set for the igroup by executing the `igroup show -v`

Note: For more information on ALUA, see the *Data ONTAP Block Access Management Guide for FCP and iSCSI*. In particular, see the section “Enabling ALUA for a Fibre Channel igroup.”

The following command line displays information about the `cfmode` on the storage system and shows that ALUA is enabled. (To make the information on ALUA easier to locate, it is shown in bold.)

```
filerA# igroup show -v
Tester (FCP):
  OS Type: solaris
  Member: 10:00:00:00:c9:4b:e3:42 (logged in on: 0c)
```

```
Member: 10:00:00:00:c9:4b:e3:43 (logged in on: vtic)
ALUA: Yes
```

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 instructions that shipped with your hardware.

You should have completed the following tasks:

- Installed the switch in a system cabinet or rack.
- Confirmed that the Host Utilities support this switch.
- Turned on power to the switch.
- Configured the network parameters for the switch, including its serial parameters and IP address.

Related information

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

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 provides the following advantages:

- It simplifies configuring the host.
- It makes information 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.

NetApp recommends that your configuration have no more than eight paths per LUN. For more information about zoning, see the NetApp Interoperability Matrix.

For more information on supported switch topologies, see the iSCSI and Fibre Channel Configuration Guide.

Related information

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

iSCSI and Fibre Channel Configuration Guide -

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

Power up equipment in the correct order

The different pieces of hardware communicate with each other in a prescribed order, which means that problems occur if you turn on power to the equipment in the wrong order.

Use the following order 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:

- The storage system's console
- A remote login that you established from the host

Possible iSCSI issues

The following sections describe possible issues that might occur when you are working with iSCSI.

Next topics

[\(iSCSI\) Verify the type of discovery being used](#) on page 139

[\(iSCSI\) Bidirectional CHAP doesn't work](#) on page 140

[\(iSCSI\) LUNs are not visible on the host](#) on page 140

(iSCSI) Verify the type of discovery being used

The iSCSI version of the Host Utilities supports iSNS, dynamic (SendTarget) and static discovery.

You can use the `iscsiadm` command to determine which type of discovery you have enabled.

This example uses the `iscsiadm` command to determine that dynamic discovery is being used.

```
$ iscsiadm list discovery
Discovery:
  Static: disabled
```

```
Send Targets: enabled
iSNS: disabled
```

(iSCSI) Bidirectional CHAP doesn't work

When you configure bidirectional CHAP, make sure you supply different passwords for the `inpassword` value and the `outpassword` value.

If you use the same value for both of these passwords, CHAP appears to be set up correctly, but it does not work.

(iSCSI) LUNs are not visible on the host

Storage system LUNs are not listed by the `iscsiadm list target -S` command or by the `sanlun lun show all` command.

If you encounter this problem, verify the following configuration settings:

- Network connectivity—Verify that there is TCP/IP connectivity between the host and the storage system by performing the following tasks:
 - From the storage system command line, ping the host.
 - From the host command line, ping the storage system.
- Cabling—Verify the cables between the host and the storage system are properly connected.
- System requirements—Verify that you have the correct Solaris operating system (OS) version, version of Data ONTAP, and other system requirements. See the NetApp Interoperability Matrix.
- Jumbo frames—If you are using jumbo frames in your configuration, ensure that jumbo frames are enabled on all devices in the network path: the host Ethernet NIC, the storage system, and any switches.
- iSCSI service status—Verify that the iSCSI service is licensed and started on the storage system. For more information about licensing iSCSI on the storage system, see the Data ONTAP Block Access Management Guide for your version of Data ONTAP.
- Initiator login—Verify that the initiator is logged in to the storage system by entering the `iscsi initiator show` command on the storage system console.

If the initiator is configured and logged in to the storage system, the storage system console displays the initiator node name and the target portal group to which it is connected.

If the command output shows that no initiators are logged in, verify that the initiator is configured according to the procedure described in the section on “Configuring the initiator.”
- iSCSI node names—Verify that you are using the correct initiator node names in the `igroup` configuration.

On the storage system, use the `igroup show -v` command to display the node name of the initiator. This node name must match the initiator node name listed by the `iscsiadm list initiator -node` command on the host.

- LUN mappings—Verify that the LUNs are mapped to an igroup that also contains the host. On the storage system, use one of the following commands:
 - `lun show -m` (displays all LUNs and the igroups they are mapped to)
 - `lun show -g igroup_name` (displays the LUNs mapped to the specified igroup)
- If you are using CHAP, verify that the CHAP settings on the storage system and host match. The incoming user names and password on the host are the outgoing values on the storage system. The outgoing user names and password on the storage system are the incoming values on the host. For bidirectional CHAP, the storage system CHAP username must be set to the storage system's iSCSI target node name.

Related information

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

Diagnostic utilities for the Host Utilities

The Host Utilities provide several command-line utilities that provide diagnostic information you can use for troubleshooting issues with your system, when requested to do so by Technical Support.

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 Technical Support.

The utilities are:

- `controller_info`. The `controller_info` utility collects information about your storage system.

Note: Previous versions of the Host Utilities used the `filer_info` utility to collect storage system information. You can still use `filer_info`; however, it uses RSH. Data ONTAP 8.0 and later use ZAPI, which means that `controller_info` works without requiring any additional configuration. To use `filer_info` with Data ONTAP 8.0, you must enable RSH.
- `solaris_info`. The `solaris_info` utility collects information about your host.
- `<switch>_info`. The `<switch>_info` utility collects information about your switch configuration. There is a utility for each switch: `brocade_info`, `cisco_info`, `mcdata_info`, and `qlogic_info`.

Note: IPv6 addressing is not supported by the diagnostic utilities. For information on how to collect this type of information manually, see the KnowledgeBase article 43800.

Next topics

Executing the diagnostic utilities (general steps) on page 142

Collecting host diagnostic information on page 142

Collecting storage system diagnostic information on page 143

Collecting switch diagnostic information on page 144

Executing the diagnostic utilities (general steps)

Although some of the diagnostic utilities have additional arguments, the general steps for executing these utilities are the same.

Steps

1. On the host console, change to the directory containing the Host Utilities' SAN Toolkit: `cd /opt/NTAP/SANToolkit/bin`
2. Enter the name of the diagnostic utility you are executing and the arguments for that utility.

The arguments are listed both in the specific section for the diagnostic utility and the man page for that utility.

The utility gathers the information and places it in a compressed file that you can send to Technical Support for analysis.

Collecting host diagnostic information

If you are having a problem, Technical Support might ask you to run the Host Utilities' `solaris_info` utility to collect diagnostic information about your Solaris host configuration.

Before you begin

Note: IPv6 addressing is not supported by the SAN Toolkit diagnostic utilities. For more information on how to collect this information manually, see the Release Notes.

Steps

1. On the host, change to the directory containing the SAN Toolkit utilities: `cd /opt/NTAP/SANToolkit/bin`
2. At the host command prompt, enter the following command:

```
solaris_info [-d path] [-n file_name] [-f]
```

The options have the following values:

- The `-d path` option lets you to specify a specific path in which to place the resulting output file. The default path is `/tmp/NTAP`.
- The `-n file_name` option specifies the name of the output file that the utility produces.
- The `-f` option forces the utility to overwrite any existing output file that is in the same directory and has the same name as the one being created. You might have an existing file if you previously ran the utility. If you do not use the `-f` option and you have an existing output file in that directory, the utility tells you to delete the file or rename it.

When you execute this command, it creates a compressed file containing the output of several host status commands and places this file in the directory you specify. It also populates the directory with files containing output from numerous commands run on the Solaris system and the contents of the `/var/adm/message` file. If you do not specify a directory, the utility places the file in the default directory. The name of the file is `file_name.tar.Z`.

3. Send the file produced by this utility to Technical Support for analysis.

Collecting storage system diagnostic information

If you are having a problem, Technical Support might ask you to run the Host Utilities' `controller_info` or `filer_info` utility to collect diagnostic information about your storage system configuration.

Before you begin

You can use the `controller_info` utility with Data ONTAP 7.2 and later. The `controller_info` utility uses the ZAPI interface while the `filer_info` utility uses RSH. Data ONTAP 8.0 and later use the ZAPI interface as the default so `controller_info` works without requiring additional setup steps.

Note: To use `filer_info` with Data ONTAP 8.0, you must enable RSH. Otherwise, you will receive an error message stating that there was a failure to connect.

These utilities use the same login and password values to access both controllers in an active/active storage system configuration. The controllers must be configured with identical login and password values for the utility to work correctly.

Both `controller_info` and `filer_info` collect the similar information, though `controller_info` collects more information than `filer_info` does. The `controller_info` utility is being added with this release and will be the primary storage system diagnostic utility going forward. The examples below use `controller_info`. Unless otherwise specified, all the information is the same for `filer_info`.

Note: IPv6 addressing is not supported by the SAN Toolkit diagnostic utilities. For more information on how to collect this information manually, see the Release Notes.

Steps

1. On the host, change to the directory containing the SAN Toolkit utilities: `cd /opt/NTAP/SANToolkit/bin`
2. At the host command prompt, enter the following command:

```
controller_info [-s][-d path] [-l password] -n controller_name [-u command] [-f]
```

The options have the following values:

- The `-s` option lets you specify whether a secure connection (SSL) should be used to communicate with the storage system(s).

Note: You cannot use the `-s` option with the `filer_info` utility. It is only supported with the `controller_info` utility.

- The `-d path` option lets you specify a specific path where the utility will place the resulting output file. The default path is `/tmp/netapp/`.
- The `-l password` option lets you specify the password needed to access the storage system.
- The `-n controller_name` option specifies the target of the command. This is the name or the IP address of the storage system being accessed by the utility.
- The `-u command` option lets you enter a command that the utility will execute.
- The `-f` option forces the utility to overwrite any existing output file that is in the same directory and has the same name as the one being created. You might have an existing file if you previously ran the utility. If you do not use the `-f` option and you have an existing output file in that directory, the utility tells you to delete the file or rename it.
- The `-h` option displays the help information and exits.

When you execute this command, it creates a compressed file containing the output of several storage system status commands and places this file in the directory you specify. If you do not specify a directory, it places the file in the default directory. The name of the file is `storagesystem_controllername.tar.Z`.

3. Send the file produced by this utility to Technical Support for analysis.

Collecting switch diagnostic information

If you are having a problem, a Technical Support representative might ask you to run the Host Utilities' switch diagnostic utility for your switch: `brocade_info`, `cisco_info`, `mcddata_info`, or `qlogic_info`.

Before you begin

Note: IPv6 addressing is not supported by the SAN Toolkit diagnostic utilities. For more information on how to collect this information manually, see the Release Notes.

Steps

1. On the host, change to the directory containing the switch utilities:

```
cd /opt/NTAP/SANToolkit/bin
```

2. At the command prompt, enter the command:

```
switch_brand_info [-d path] [-l user:password] [-l images_ftp_password]  
-n switchname -f
```

The options have the following values:

- The `switch_brand` is the brand name for the switch. The Host Utilities provide `brocade_info`, `cisco_info`, `mcddata_info`, and `qlogic_info`.

- The `-d path` option enables you to specify a path and filename for the output. The default path is `/tmp/netapp`.
- The `-l user:password` option specifies the user name and password required to access the switch.

Note: If you are using a Brocade switch and do not specify this option, the utility uses Brocade's default settings of `admin:password`.

- The `-l images_ftp_password:password` option specifies the user name and password required to access the switch.

Note: This option applies only to the `qlogic_info` utility. You cannot use it with the other switch diagnostic utilities.

- The `-n switchname` option specifies the name or IP address of the switch from which data is collected.
- The `-f` option forces the utility to overwrite any existing output file that is in the same directory and has the same name as the one being created. You might have an existing file if you previously ran the utility. If you do not use the `-f` option and you have an existing output file in that directory, the utility tells you to delete the file or rename it.
- The `-h` option displays the help information and exits.

When you execute this command, it creates a compressed file containing the output of several switch status commands and places this file in the directory you specify. If you do not specify a directory, it places the file in the default directory. The name of the file is `switch_brand_switch_switchname.tar.Z`.

The following example uses the `brocade_info` utility. The utility uses the user “chris” and the password “smith” to connect to the switch, which is named `brcd1`.

```
host1> brocade_info -d /opt/NTAP/SANToolkit/bin/ -l chris:smith -n brcd1
```

The information from this switch is placed in the `brocade_switch_brcd1.tar.Z` file.

3. Send the file produced by this utility to Technical Support for analysis.

Possible MPxIO issues

The following sections describe issues that can occur when you are using the Host Utilities in an MPxIO environment.

Next topics

[\(MPxIO\) sanlun does not show all adapters](#) on page 146

[\(MPxIO\) Solaris log message says data not standards compliant](#) on page 146

(MPxIO) sanlun does not show all adapters

In some cases, the `sanlun lun show all` command does not display all the adapters. You can display them using either the `luxadm display` command or the `mpathadm` command.

When you use MPxIO, there are multiple paths. MPxIO controls the path over which the `sanlun SCSI` commands are sent and it uses the first one it finds. This means that the adapter name can vary each time you issue the `sanlun lun show` command.

If you want to display information on all the adapters, use either the `luxadm display` command or the `mpathadm` command. For the `luxadm display` command, you would enter

```
luxadm display -vdevice_name
```

Where *device_name* is the name of the device you are checking.

(MPxIO) Solaris log message says data not standards compliant

When running the Host Utilities, you may see a message in the Solaris log saying that data is not standards compliant. This message is the result of a Solaris bug.

```
WARNING: Page83 data not standards compliant
```

This Solaris log message is the result of a Solaris bug and has been reported to Sun. The Solaris initiator implements an older version of the SCSI Spec.

The NetApp SCSI target is standards compliant, so ignore this message.

(Veritas DMP/LPFC) The Solaris Host Utilities create_binding program

The Solaris Host Utilities provide the create_binding program as a tool to help Veritas DMP environments set up the necessary bindings. The following sections provide information on working with the create_binding program.

Next topics

[\(Veritas DMP/LPFC\) The create_binding program](#) on page 147

[The rules create_binding uses for WWPN bindings](#) on page 148

[\(Veritas DMP/LPFC\) Prerequisites for using create_binding](#) on page 148

[\(Veritas DMP/LPFC\) Storage system prerequisites for using create_binding](#) on page 149

[\(Veritas DMP/LPFC\) Creating WWPN bindings with create_binding](#) on page 149

(Veritas DMP/LPFC) The create_binding program

The create_binding program is a tool provided by the Host Utilities 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.

Note: You only need to create persistent binding if you are using a host running the Veritas DMP stack with LPFC. If you are using the MPxIO stack, you do not need to create persistent bindings.

When you upgrade an installation that uses WWNN bindings to one that uses WWPN bindings, the create_binding program saves the existing /kernel/drv/lpfc.conf file to /kernel/drv/lpfc.conf.NTAP.[date-stamp]. The program 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 that do not have existing WWPN bindings, the create_binding program accesses the storage system's WWPN list and automatically binds the WWPNs to target IDs and updates the /kernel/drv/lpfc.conf file.

Note: You can also create bindings using the hbacmd command or HBAnyware. For more information on using hbacmd, see the sections on it in this document. For information on using HBAnyware, see the documentation that comes with that tool.

The rules `create_binding` uses for WWPN bindings

The `create_binding` program uses certain rules for creating WWPN bindings.

They are:

- Initiators bind to accessible targets only.
- For multipath configurations with Veritas VxVM and DMP software, the program 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 program 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 program 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 (`cfmode`) is set to single-image, the program creates bindings for all available target ports.
- The program 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.

Note: 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.

(Veritas DMP/LPFC) Prerequisites for using `create_binding`

To use the `create_binding` program, your configuration must meet certain prerequisites.

- The LPFC driver must be installed.
- The FC Solaris Host Utilities must be installed.
The `create_binding` program is bundled with the Solaris Host Utilities software. You must install this kit to obtain the `create_binding` software.
- Switch cabling must be connected.
The `create_binding` program checks the availability of paths from the host to the storage system. All cables must be connected correctly between the host and the storage system.
- Zoning must be configured

If your system uses zoning, 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` program 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 program.

(Veritas DMP/LPFC) Storage system prerequisites for using create_binding

To use the `create_binding` program, your storage system must meet certain prerequisites.

- The FC protocol service must be licensed and running 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>`

(Veritas DMP/LPFC) Creating WWPN bindings with create_binding

To execute the `create_binding` program, you must perform several steps.

Steps

1. Ensure that you are logged in to the host and have root privileges.
2. Change to the `/opt/NTAP/SANToolkit/bin` directory:

```
cd /opt/NTAP/SANToolkit/bin
```

3. Enter the following command:

```
create_binding -l 'login:passwd,login:passwd' -n filer [-d | -D]
```

- `l login:passwd` 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 storage_system` is the storage system's host name.

For cluster configurations, you specify one member of the cluster. The program detects and configures bindings for the partner storage system.

- `d` enables the program to process information for Veritas DMP configurations.

Use the `-d` option if you have already installed VxVM with DMP. If you use the `-d` option for Veritas DMP configurations, the program automatically configures WWPN bindings for all accessible paths to the storage system.

- `D` enables the program 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.

The following is an example of a `create_binding` command line.

```
./create_binding -l root -n storage_system_1 -d
```

4. Take one of the actions below:

- If you have a single-attached or direct-attached configuration, the program 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 `create_binding` program in Step 3, the program displays a warning. The `create_binding` program will not discover multiple paths to the storage system, so you must take the following actions:
 - Enter `n` to exit the program.
 - Rerun the `create_binding` program with the `-d` option.
 - Enter `y` at the prompt to continue the program and bind by WWPN.
The program displays the new WWPN bindings and prompts you to save them to the `/kernel/drv/lpfc.conf` file.
Enter `y` to save the new WWPN bindings and proceed to Step 5.
- If you have VxVM and DMP installed without a multipathing configuration and you did not use the `-d` option of the program in Step 3, the program 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 `/kernel/drv/lpfc.conf` file.
- Proceed to Step 5.

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:

```
name="sd" parent="lpfc" target=0 lun=0;
```

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 `/kernel/drv/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:

```
touch /reconfigure  
/sbin/init 6
```

Note: If you're performing multiple configuration steps that require rebooting, you may want to perform all those steps and then reboot.

SAN boot LUNs in a Solaris Veritas DMP environment with FC

You can set up a SAN boot LUN to work in a Veritas DMP environment that using the FC protocol and running the Solaris Host Utilities.

Note: At the time this document was produced, Veritas DMP environments only supported SAN boot LUNs with the FC protocol. The Veritas DMP environment did not support SAN boot LUNs in an iSCSI environment.

The method you use for creating a SAN boot LUN and installing a new OS image on it in a DMP environment can vary, depending on the drivers you are using

The sections that follow provide steps for configuring a SAN boot LUN and installing a new OS image on it. You can apply these steps to most configurations.

There are also other methods for creating a SAN boot LUN. This document does not describe other methods for creating bootable LUNs, such as creating 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 4.4 Configuration Guide*, which is available in PDF format on the Sun site.

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 .

Next topics

[\(Veritas DMP\) Prerequisites for creating a SAN boot LUN](#) on page 154

[\(Veritas DMP\) SAN boot configuration overview](#) on page 154

[\(Veritas DMP/native\) About setting up the Sun native HBA for SAN booting](#) on page 155

[\(Veritas DMP/LPFC\) About setting up the LPFC HBA for SAN booting](#) on page 160

[\(Veritas DMP/native\) Methods for installing directly onto a SAN boot LUN](#) on page 162

[\(Veritas DMP\) Installing the bootblk](#) on page 163

[\(Veritas DMP\) Copying the boot data](#) on page 163

[\(Veritas DMP\) Information on creating the bootable LUN](#) on page 165

[\(Veritas DMP\) Partitioning the bootable LUN to match the host device](#) on page 166

[\(Veritas DMP\) What OpenBoot is](#) on page 169

Related information

[NetApp Interoperability Matrix - http://now.netapp.com/matrix/mtx/login.do](http://now.netapp.com/matrix/mtx/login.do)

(Veritas DMP) Prerequisites for creating a SAN boot LUN

You need to have your system set up and the Host Utilities installed before you create a SAN boot LUN for a Veritas DMP environment.

Note: SAN booting is only supported in Veritas environments that use the FC protocol. It is not supported with the iSCSI protocol.

Before attempting to create a SAN boot LUN, make sure the following prerequisites are in place:

- The 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 Qlogic-branded HBAs, the FCode is available on the QLogic site.
 - For Sun-branded QLogic HBAs, the FCode is available as a patch from Sun.
- If you are using Emulex HBAs, you must have the Emulex FCA utilities with EMLXemlxu and the EMLXdrv installed.
- If you are using Emulex-branded HBAs or Sun-branded Emulex HBAs, make sure you have the current FCode. The FCode is available on the Emulex site.
- If you are using QLogic-branded HBAs, you must have the SANsurfer SCLI utility installed.
- 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`
 - By using `hbacmd`

(Veritas DMP) SAN boot configuration overview

To configure a bootable LUN in a Veritas DMP environment, you must perform several tasks.

The following is a high-level overview of the tasks to help you plan the configuration.

1. Make sure the HBA is set to the appropriate Bootcode:
2. Create 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. Select 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. Modify the 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. Reboot the system.

(Veritas DMP/native) About setting up the Sun native HBA for SAN booting

Part of configuring a bootable LUN when using Veritas DMP with Sun native drivers is setting up your HBA. To do this, you may need to shut down the system and switch the HBA mode.

The Veritas DMP environment supports two kinds of Sun native HBAs. The actions you take to set up your HBA depend on the type of HBA you have. If you have

- An Emulex HBA, you must make sure the HBA is in SFS mode.
- A QLogic HBA, you must change the HBA to enable FCode compatibility.

Note: If you have LPFC HBAs, you must perform different steps.

Next topics

[\(Veritas DMP/native\) Changing the Emulex HBA to SFS mode](#) on page 155

[\(Veritas DMP/native\) Changing the QLogic HBA to enable FCode compatibility](#) on page 158

(Veritas DMP/native) 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.

About this task

Caution: 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, those entries will become invalid.

Steps

1. At the operating system prompt, issue the `init 0` command.

```
# init 0
```

2. When the `ok` prompt appears, enter the `setenv auto-boot? false` command.

```
ok setenv auto-boot? false
```

3. Enter the `reset-all` command.

```
ok reset-all
```

4. Issue the `show-devs` command to see the current device names.

The following example uses the `show-devs` command to see if the Emulex device has been set to SFS mode. In this case, executing the command shows that the device has not been set to SFS mode because the devices (shown in **bold**) are displayed as `.../lpfc`, not `.../emlxs`. See Step 5 for information setting the devices to SFS mode.

```
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
```

5. Select the first Emulex device and set it to SFS mode using the `set-sfs-boot` command. Doing this changes the devices to `emlxs` devices.

In this example the first command, `show-devs`, shows the device name. The next command, `select`, selects the device `lpfc@0`. Then the `set-sfs-boot` command sets the SFS mode. Some of the output has been truncated to make this example easier to read. To make the SFS information easy to locate, it is in **bold**.

```
ok 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...
ok select /pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@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. ***
ok
```

6. Repeat Step 5 for each Emulex device.
7. Enter the `reset-all` command to update the devices.

In this example the `reset-all` command updates the Emulex devices with the new mode.

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

8. Issue the `show-devs` command to confirm that you have changed the mode on all the Emulex devices.

The following example uses the `show-devs` command to confirm that the Emulex devices are showing up as `emlx` devices. To continue the example shown in from Step 5, the device selected there, `/pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0`, has been changed to an `emlx` device. In a production environment, you would want to change all the devices to `emlx`.

```
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

```

9. Set the `auto-boot?` back to `true` and boot the system with a reconfiguration boot.

This example uses the `boot` command to bring the system back up.

```

ok setenv auto-boot? true
ok boot -r

```

(Veritas DMP/native) 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.

Steps

1. At the operating system prompt, issue the `init 0` command.

```
# init 0
```

2. When the `ok` prompt appears, enter the `setenv auto-boot? false` command.

```
ok setenv auto-boot? false
```

3. Enter the `reset-all` command.

```
ok reset-all
```

4. Issue the `show-devs` command to see the current device names.

The following example uses the `show-devs` command to see whether there is FCode compatibility. It has been truncated to make it easier to read.

```

ok show-devs
...
/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1
/pci@7c0/pci@0/pci@8/QLGC,qlc@0
/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1/fp@0,0
/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1/fp@0,0/disk

```

```
/pci@7c0/pci@0/pci@8/QLGC,qlc@0/fp@0,0
/pci@7c0/pci@0/pci@8/QLGC,qlc@0/fp@0,0/disk
```

5. Select the first QLogic device.

This example uses the `select` command to select the first QLogic device.

```
ok select /pci@7c0/pci@0/pci@8/QLGC,qlc@0,1
QLogic QLE2462 Host Adapter Driver(SPARC): 1.16 03/10/06
```

6. If you need to set the compatibility mode to `fcode`, execute the `set -mode` command.

The following example uses the `set -mode` command to set the compatibility mode to `fcode`.

```
ok set-mode
Current Compatibility Mode: fcode
Do you want to change it? (y/n)
Choose Compatibility Mode:
0 - fcode
1 - bios
enter: 0
Current Compatibility Mode: fcode
```

7. Execute the `set -fc-mode` command and, if needed, to set the Fcode Mode to `qlc`.

The following example uses the `set -mode` command to set the mode to `qlc`.

```
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: qlc
```

8. Repeat the previous steps to configure each QLogic device.
9. Enter the `reset -all` command to update the devices.

The following example uses the `reset -all` command to update the QLogic devices with the new mode..

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

10. Set the `auto-boot?` back to `true` and boot the system with a reconfiguration boot.

This example uses the `boot` command to bring the system back up.

```
ok setenv auto-boot? true
ok boot -r
```

(Veritas DMP/LPFC) About setting up the LPFC HBA for SAN booting

Part of configuring a bootable LUN using Emulex LPFC drivers in a Veritas DMP environment 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. The following section provides information on doing this.

Note: If you have Sun native HBAs, you must perform different steps.

(Veritas DMP/LPFC) 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.

About this task

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

Steps

1. At the operating system prompt, issue the `init 0` command.

```
# init 0
```

2. When the `ok` prompt appears, enter the `setenv auto-boot? false` command.

```
ok setenv auto-boot? false
```

3. Enter the `reset-all` command.

```
ok reset-all
```

4. Issue the `show-devs` command to see the current device names.

The following example uses the `show-devs` command to see if the Emulex device has been set to SD mode. In this case, executing the command shows that the device has not been set to SD mode because the devices (shown in **bold**) are displayed as `.../emlxs`, not `.../lpfc`. See Step 5 for information setting the devices to SD mode.

```

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

```

5. Select the first Emulex device and set it to SD mode using the `set -sd-boot` command. Doing this changes the Emulex devices to lpfc devices.

In this example the `select` command selects the Emulex device `700000/emlxs@1`. The second command, `set -sd-boot`, sets the SD mode. To make the SFS information easy to locate, it is in **bold**.

```

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. ***

```

6. Repeat Step 5 for each Emulex device.
7. Enter the `reset-all` command to update the devices.

In this example the `reset-all` command updates the Emulex devices with the new mode.

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

8. Issue the `show-devs` command to confirm that you have changed the mode on all the Emulex devices.

The following example uses the `show-devs` command to confirm that the Emulex devices are showing up as `lpfc` devices. To make those devices easy to locate, they are in **bold**.

```
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
```

9. Set the `auto-boot?` back to `true` and boot the system with a reconfiguration boot.

This example uses the `boot` command to bring the system back up.

```
ok setenv auto-boot? true
ok boot -r
```

(Veritas DMP/native) Methods for 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 Solaris operating system using:

- A CD/DVD installation
- A jumpstart installation

During the installation, select the bootable LUN you created earlier.

(Veritas DMP) Installing the bootblk

The bootblk contains startup information that the Solaris host requires. You must install the bootblk onto the raw bootable LUN.

About this task

Installing the bootblk involves:

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

Steps

1. Determine the directory in `/usr/platform` where the bootblks are located. Enter:

```
uname -i
```

In this example, the output indicates that bootblks should reside in `/usr/platform/sun4u` (shown in **bold**).

```
#uname -i
SunOS shsun17 5.10 Generic_118833-33 sun4u sparc SUNW,Sun-Fire-280R
```

2. Write the bootblk. Enter:

```
/usr/sbin/installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk
/dev/rdisk/<device_name>
```

The following example installs bootblks in the `/usr/platform/'uname -i'` directory on `/dev/rdisk/c2t500A098786F7CED2d11s0`.

```
#!/usr/sbin/installboot /usr/platform/sun4u/lib/fs/ufs/bootblk
/dev/rdisk/c2t500A098786F7CED2d11s0
```

(Veritas DMP) Copying the boot data

Copying the boot data to the bootable LUN involves several steps.

About this task

You must:

- Create and mount 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.
- Move boot data from the host device onto the bootable LUN.
- Editing the `/etc/vfstab` file to reference the bootable LUN.

Steps

1. Create a file system on the bootable LUN that matches the file system in the `/etc/vfstab` file of the source device. Enter:

```
newfs <device>
```

The following example creates a UFS file system at `/dev/rdisk/c2t500A098786F7CED2d11s0`.

```
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
The next example shows the corresponding entry (in bold) for the file
system in the vfstab file.
#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 -
```

2. Mount the root directory of the bootable LUN. Enter:

```
mount <filesystem_name><mountpoint>
```

The following example mounts the file system on the bootable LUN under `/mnt`.

```
#mount /dev/dsk/c2t500A098786F7CED2d11s0 /mnt
```

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:

```
ufsdump 0f - <source_boot_device> | (cd /<mntpoint_of_bootable_lun>;
ufsrestore rf -)
```

The following example 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.

```
#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/rdisk/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
```

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 the `sanlun lun show -p` command. The `sanlun` command comes with the Host Utilities.

The following shows the `/etc/vfstab` entry for the bootable LUN (bold).

```
#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/rdisk/c2t500A098786F7CED2d11s0

/              ufs          1          yes -
/devices        -          /devices    devfs        -          no          -
ctfs            -          /system/contract ctfs        -          no          -
objfs           -          /system/object objfs        -          no          -
swap            -          /tmp        tmpfs        -          yes         -
```

(Veritas DMP) Information on creating the bootable LUN

After setting up the HBAs, you must create the LUN you want to use a boot LUN.

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

In addition, you must partition the bootable LUN so that it matches the partitions on the host boot device. Partitioning the LUN involves:

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

(Veritas DMP) Partitioning 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.

Steps

1. Select the current Solaris boot device:

- a. Execute the Solaris format utility. Enter the `/usr/sbin/format` utility.

The system displays Available Disk Selections.

In this example, the format utility displays information about the disks on your system.

```
# 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
```

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.

- b. 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.

The system selects the Solaris boot disk and displays partition options.

2. Look in the `/etc/vfstab` file and write down the information about the boot file systems.

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.

```
# 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/rdisk/c1t1d0s0 / ufs 1 no -
/devices - /devices devfs - no -
ctfs - /system/contract ctfs - no -
objfs - /system/object objfs - no -
swap - /tmp tmpfs - yes -
```

3. Display partition information for the current Solaris boot disk:

- a. At the `format>` prompt, enter

p

to display partition menu options.

- b. At the `partition>` 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 `/etc/vfstab` file.

The system displays the partitions for the boot device.

In this example shows 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)

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

```

c. At the `partition>` prompt, enter

q

to go back to the top level `format>` prompt.

d.

4. 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.
5. Use the `format` utility to create the root same size or larger and swap slice for the SAN boot disk.
 - a. At the prompt, type

disk

to get a list of disks.

- b. At the `specify disk` prompt, enter the number that corresponds to the LUN mapped to host controller and target0.

Make a note of the identifying information for the LUN you select. You will use this information when you bind the bootable LUN to the adapter and create a boot alias as part of the task to modify OpenBoot. Pay close attention to the “cXtYdZ” information.

This example selects c2t500A098786F7CED2d11 as the bootable LUN.

```
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]

6. Display partition information for the current bootable LUN:

a. At the `format>` prompt, enter

p

to display partition menu options.

b. At the `partition>` prompt, enter

p

to print the partition table for the current Solaris boot device.

The system displays the partitions for the bootable LUN.

7. 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.

8. At the `partition>` prompt, enter

label

to write the new label to the disk.

9. Exit the `format` utility.

a. At the `partition>` prompt, enter

q

b. At the `format>` prompt, enter

q

(Veritas DMP) What OpenBoot is

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.

When you're setting up a SAN boot LUN, you need to modify OpenBoot to create an alias for the bootable LUN. The alias substitutes for the device address during subsequent boot operations.

The steps you need to perform to modify OpenBoot differ based on whether you are using Sun native drivers or LPFC drivers.

The following sections provide information on working with OpenBoot.

Next topics

[\(Veritas DMP/native\) Modifying OpenBoot with Sun native drivers](#) on page 170

[\(Veritas DMP/LPFC\) Modifying OpenBoot for the Emulex LPFC driver](#) on page 172

(Veritas DMP/native) Modifying OpenBoot with Sun native drivers

Modifying OpenBoot when you are using Sun native drivers 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.

Steps

1. Activate the OpenBoot environment by performing the following steps:

- a. Halt the Solaris operating system by entering the command:

```
init 0
```

The Open Boot program displays the OpenBoot prompt.

- b. Stop automatic reboot during OpenBoot configuration by entering the command:

```
setenv auto-boot? false
```

- c. Reset the host and activate your changes by entering:

```
reset-all
```

2. Run the show-disks command to get the correct device path name for your FC 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.

In this example, the output shown here continues the example used in the section on *Partitioning the bootable LUN to match the host device*. It uses the information that was written down when you executed the format command. The LUN is “w500a098786f7ced2,b” and the device path is “/pci@1d,700000/emlx@2/fp@0,0/ssd.”

```
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
```

The following example shows the type of output you see when you execute the show-disks command.

```
{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:
```

3. Create an alias for the bootable LUN. Enter:

```
nvalias <alias_name><boot_LUN_pathname>
```

This example creates an alias called “sanbootdisk”.

```
{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
```

4. To allow the system to autoboot to the LUN, add the LUN’s alias to your boot device list. Enter the following command:

```
setenv boot-device <alias_name>
```

Note: Before adding the new alias, use the `printenv boot-device` 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 `setenv boot-device` command line after the bootable LUN’s alias. That alias must always be the first argument. The format for the command would be

```
setenv boot-device <bootable_lun_alias_name> <existing_alias1>
<existing_alias2> ...
```

This example assumes that you’ve created the alias `sanbootdisk` for your bootable LUN. You check the boot device list and discover that there are already two boot devices in the list: `disk` and `net`. To add the alias `sanbootdisk`, you would enter the `setenv boot-device` command followed by `sanbootdisk` and the other two aliases in the list.

```
{0} ok setenv boot-device sanbootdisk disk net
```

5. Enable auto-boot. Enter:

```
setenv auto-boot? true
```

6. Boot the LUN. Enter:

```
boot <alias_name>
```

This example boots the alias for `/pci@1d,700000/emlx@2/fp@0,0/disk@w500a098786f7ced2,b`

```
{0} ok boot sanbootdisk
```

7. 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.

(Veritas DMP/LPFC) Modifying OpenBoot for the Emulex LPFC driver

Modifying OpenBoot for LPFC involves several steps.

About this task

You must:

- Run `san lun` to get information about your LUN.
- Map the HBA instance to a physical path.
- Bind the target and the LUN to the HBA.
- Create 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
```

Steps

1. Run the `san lun show -p` 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 `/etc/vfstab` file on the SAN boot LUN.

Note: The `san lun` command is in the `/opt/NTAP/SANToolkit/bin/` directory.

This example uses the `san lun show -p` command to display information about the LUN.

```
# /opt/NTAP/SANToolkit/bin/san 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
state     type          device          host      primary partner
          type          filename        HBA      filer   filer
          type          filename        HBA      port   port
-----
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

2. Get the device path for the lpfc instance to the device path mapping from the /etc/path_to_inst file.

This example shows how you can use the grep command to get the device path.

```
# grep \"lpfc\" /etc/path_to_inst

\"/pci@1d,700000/lpfc@2\" 0 \"lpfc\"
\"/pci@1d,700000/lpfc@2,1\" 1 \"lpfc\"
```

3. Document the WWPN binding for your controllers. The persistent binding information can be found in the /kernel/drv/lpfc.conf file. Look for the variable named 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.

This example shows the binding for the sample configuration.

```
fcp-bind-wwpn=\"500a098486c7ced2:lpfc0t0\",
               \"500a098386c7ced2:lpfc0t1\",
               \"500a098286c7ced2:lpfc1t0\",
               \"500a098186c7ced2:lpfc1t1\",
               \"500a098496c7ced2:lpfc0t2\",
               \"500a098396c7ced2:lpfc0t3\",
               \"500a098296c7ced2:lpfc1t2\",
               \"500a098196c7ced2:lpfc1t3\";
```

4. Activate the OpenBoot environment by performing the following steps:
 - a. Halt the Solaris operating system by entering the command:


```
init 0
```

The Open Boot program displays the OpenBoot prompt.
 - b. Stop automatic reboot during OpenBoot configuration by entering the command:


```
setenv auto-boot? false
```
 - c. Reset the host and activate your changes by entering:


```
reset-all
```
5. Bind the HBA to the LUN and the target.

This step uses all of the information gathered in steps 1 through 3. To determine the value of target-in-hex, you need to refer to the persistent binding entries you documented in step 3. The target is the value found in position Y of wwpn: lpfcXtY. Convert this number to hex and use it in the setup below.

```
select /device_path_to_card
wwpn <wwpn-of-target> <lun-in-hex> <target-in-hex> set-boot-id
reset-all
```

The following example show the type of output that occurs when you perform this task.

```
ok select /pci@1d,700000/lpfc@2
ok wwpn 500a098486c7ced2 f5 0 set-boot-id
Flash data structure updated.
Signature 4e45504f
Valid_flag 14e
Host_did 0
Enable_flag 5
SFS_Support 0
Topology_flag 0
Link_Speed_flag 0
Diag_Switch 0
Boot_id 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
*** Type reset-all to update. ***
ok reset-all
```

6. Run the show-disks command to get the correct device path name for your FC 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 “sd” as the terminating characters and those characters may differ from what was seen in the format output.

In this example, the device uses "sd" to represent the disk device.

```
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/lpfc@2,1/sd
f) /pci@1d,700000/lpfc@2/sd
q) NO SELECTION
Enter Selection, q to quit:
Enter Selection, q to quit:
```

7. Create an alias for the bootable LUN. Enter:

```
nvalias <alias_name><boot_LUN_pathname>
```

This example creates an alias called “sanbootdisk”.

```
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
```

8. To allow the system to autoboot to the LUN, add the LUN's alias to your boot device list. Enter the following command:

```
setenv boot-device <alias_name>
```

Note: Before adding the new alias, use the `printenv boot-device` 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 `setenv boot-device` command line after the bootable LUN's alias. That alias must always be the first argument. The format for the command would be:

```
setenv boot-device <bootable_lun_alias_name> <existing_alias1>  
<existing_alias2> ...
```

This example assumes that you've created the alias `sanbootdisk` for your bootable LUN. You check the boot device list and discover that there are already two boot devices in the list: `disk` and `net`. To add the alias `sanbootdisk`, you would enter the `setenv boot-device` command followed by `sanbootdisk` and the other two aliases in the list.

```
{0} ok setenv boot-device sanbootdisk disk net
```

9. Enable auto-boot. Enter:

```
setenv auto-boot? true
```

10. Boot the LUN. Enter:

```
boot <alias_name>
```

This example boots the alias for `/pci@1d,700000/lpfc@2/sd@0,f5`

```
{0} ok boot sanbootdisk
```

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.

SAN boot LUNs in an MPxIO environment

You can set up a SAN boot LUN to work in an MPxIO environment that is running the Solaris Host Utilities.

The method you use for creating a SAN boot LUN and installing a new OS image on it in a MPxIO environment can vary, depending on the drivers you are using.

The sections that follow provide steps for configuring a SAN boot LUN and installing a new OS image on it. The examples in these sections use a Solaris host to configure a SAN boot LUN and install a new OS image onto it.

The sections that follow provide examples of copying data from an existing, locally booted server using:

- A direct installation
- A file system dump and restore for a SPARC system with MPxIO disabled
- A file system dump and restore for an x86/x64 system with MPxIO disabled
- A file system dump and restore for an x86/x64 system with MPxIO enabled

You can use these procedures to set up SAN boot LUNs for most MPxIO configurations.

There are also other methods for creating a SAN boot LUN. This document does not describe other methods for creating bootable LUNs, such as creating configurations that boot multiple volumes or use diskless servers.

Refer to Solaris documentation for details about additional configuration methods. For complete details on the procedures that follow, see Sun's document *Sun StorEdge SAN Foundation Software 4.4 Configuration Guide*, which is available in PDF format on the Sun site at docs.sun.com. Make sure that you use the FCode specified in the NetApp Interoperability Matrix .

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 .

Next topics

[\(MpxIO\) Prerequisites for creating a SAN boot LUN](#) on page 178

[\(MPxIO\) Options for setting up SAN booting](#) on page 178

[\(MPxIO\) Performing a direct install to create a SAN boot LUN](#) on page 178

[SPARC systems without MPxIO: Copying data from locally booted server](#) on page 179

[x86/x64 with MPxIO systems: Copying data from a locally booted disk](#) on page 182

[x86/x64 without MPxIO systems: Copying data from locally booted server](#) on page 185

Related information

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

(MPxIO) Prerequisites for creating a SAN boot LUN

You need to have your system set up and the Host Utilities installed before you create a SAN boot LUN in an MPxIO environment.

Note: SAN booting is only supported in MPxIO environments that use the FC protocol. It is not supported with the iSCSI protocol.

Before attempting to create a SAN boot LUN, make sure the following prerequisites are in place:

- The 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 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.

(MPxIO) Options for setting up SAN booting

There are two ways to set up SAN booting in an MPxIO environment.

- Copy the data from a locally booted device to a SAN device. Then reconfigure the host to boot to the SAN device. On x86/x64 systems, using this method means you do not have to re-set up a currently configured host and its MPxIO settings.
- Installing the operating system directly on the SAN device. Once you do this, you must configure the SAN boot LUN on the storage system and map it to the host.

(MPxIO) Performing a direct install to create a SAN boot LUN

You can install the operating system directly onto the SAN device.

About this task

To perform a direct install, run the Solaris Installer and choose the SAN device as your install device.

Steps

1. Run the Solaris Installer.
2. Choose SAN device as your install device.

3. Now configure the SAN boot LUN on the storage system and map it to the host.

SPARC systems without MPxIO: Copying data from locally booted server

If you have a SPARC system that has MPxIO disabled, you can create a SAN boot LUN by copying data from a locally booted server using a file system dump and restore.

Steps

1. If MPxIO is enabled, disable it:
 - For Solaris 10 update 5 and later, use this `stmsboot` command line:
`stmsboot -D fp -d`
 - For Solaris 10 releases prior to update 5, use this `stmsboot` command line:
`stmsboot -d`

2. Use the `mount` command to identify the current boot device.

This example mounts the current boot device.

```
# mount
/ on /dev/dsk/c0t0d0s0
read/write/setuid/devices/intr/largefiles/logging/xattr/onerror=panic/dev=800010
on Tue Jan  2 07:44:48 2007
```

3. Use the `format` command to choose the LUN. Write down the device path for the LUN so you can refer to it later. You need to configure the LUN for
 - The root file system
 - The file systems that will be copied to the bootable LUN. You may want to create slices and place the file systems on the slices.

In addition, you might want to set up the swap slice and slices for the metadatabase replicas.

In this example, you have already run the `format` command and, based on its output, determined that you want device 4, which has device path
/pci@1d,700000/emlx@2,1/fp@0,0/ssd@w500a09838350b481,f4.

```
c1t500a09838350b481d244 <NETAPP-LUN-0.2 cyl 6142 alt 2 hd 16 sec 256>
```

4. Install the bootblock on the new bootable LUN by using the `installboot` command:

```
installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk /dev/rdisk/<boot lun>
```

In this example, the following command line installs the bootlock on the boot LUN
c1t500A09838350B481d32s0.

```
installboot /usr/platform/'uname -i'/lib/fs/ufs/bootblk
/dev/rdisk/c1t500A09838350B481d32s0
```

5. Create the SAN file system that you will dump the current file system to.

In this example, the following command line creates a new file system.

```
newfs /dev/dsk/c1t500A09838350B481d32s0
```

6. Mount the file system that you will use when you copy the boot data.

In this example, the following example mounts the SAN file system.

```
mount /dev/dsk/c1t500A09838350B481d32s0 /mnt/boot lun
```

7. Create the required directory structure on the bootable LUN and copy the boot data. Enter:

```
ufsdump 0f - <source_boot_device> | cd /<mntpoint of bootable_ lun>;
ufsrestore rf -)
```

Note: If your configuration boots off more than one device, you must create and configure a bootable LUN that matches each boot device on the host. If you are copying non-root boot partitions, see the Sun's document *Sun StorEdge SAN Foundation Software 4.4 Configuration Guide* for instructions.

The following example copies the information from c0t0d0s0.

```
# ufsdump 0f - /dev/rdisk/c0t0d0s0 | (cd /mnt/boot lun; ufsrestore rf -)
```

8. Edit the /etc/vfstab file to on the SAN file system. Change the swap, root, and file systems in the file to show the boot device instead of the local device. One way to do this is to edit the file using vi:

```
cd /mnt/boot lun/etc/
```

```
vi vfstab
```

The following example shows the vfstab entry for the bootable LUN.

#device	device	mount	FS	fsck	mount
#to mount	to fsck	point	type	pass	at boot
options					
#					
fd -	/dev/fd fd	-	no	-	
/proc -	/proc proc	-	no	-	
/dev/dsk/c1t500A09838350B481d32s1		-	-	swap	-
no -					
/dev/dsk/c1t500A09838350B481d32s0		/dev/rdisk/c1t500A09838350B481d32s0			
/	ufs 1	no	-		
/dev/dsk/c1t500A09838350B481d32s6		/dev/rdisk/c1t500A09838350B481d32s6			
/globaldevices	ufs 2	yes	-		
/devices	-	/devices	devfs	-	no
ctfs -	/system/contract	ctfs	-	no	-

objfs	-	/system/object	objfs	-	no	-
swap	-	/tmp	tmpfs	-	yes	-

9. Umount the file system on the bootable LUN.

```
umount /mnt/boot lun
```

10. Run the show-disks command to get the correct device path name for your FC 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.

In this example, the output shows that the device uses "disk" to represent the disk device:

```
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:
```

Using the device name you wrote down in Step 3, boot to the bootable LUN. Keep two things in mind as you do this:

- You can specify the slice of the bootable LUN by using ":a" for slice zero, ":b" for slice 1, and so on.
- If you are using Emulex HBAs or StorageTek HBAs based on Emulex technology, replace the "ssd" part of the string with "disk".

The following examples show command lines for QLogic HBAs and Emulex HBAs.

QLogic:

```
# boot /pci@1f,2000/QLGC,qlc@1/fp@0,0/ssd@w500a09838350b481,20:a
```

Emulex:

```
# boot /pci@1d,700000/SUNW,emlxs@2/fp@0,0/disk@w500a09818350b481,20:a
```

11. Using the device name you wrote down in Step 3, set the boot-device environment variable and boot the system.

Note: You can specify the slice of the bootable LUN by using ":a" for slice zero, ":b" for slice 1, and so on.

This example sets the boot device and reboots the system.

```
ok setenv boot-device
/pci@1d,700000/emlx@2,1/fp@0,0/disk@w500a09838350b481,f4:a
ok boot -r
```

12. Re-enable MPxIO to provide path redundancy.

- For Solaris 10 update 5 and later, use this stmsboot command line:
stmsboot -D fp -e
- For Solaris 10 releases prior to update 5, use this stmsboot command line:
stmsboot -e

13. Verify that you booted to the correct LUN by using the df -k command.

In addition, configure the system crash dump for the new bootable LUN boot device using the dumpadm -d command:

```
dumpadm -d /dev/rdisk/c1t500A09838350B481d32s1
```

x86/x64 with MPxIO systems: Copying data from a locally booted disk

If you have an x86/x64 system that is using MPxIO, you can copy data from a locally booted server.

Steps

1. Use the mount command to identify the current boot device.

The following example mounts the current boot device.

```
# mount / on /dev/dsk/c1d0s0
read/write/setuid/devices/intr/largefiles/logging/xattr/onerror=panic/dev=1980000
on Tue Jan 2 12:50:22 2007
```

2. Use the format command to choose the LUN. Write down the device path for the LUN so you can refer to it later. You need to configure the LUN for

- The root file system
- The file systems that will be copied to the bootable LUN. You may want to create slices and place the file systems on the slices.

In addition, you might want to set up the swap slice and slices for the metadata database replicas.

In this example, you have already run the format command and, based on its output, determined that you want number 33, which has the device path

```
/scsi_vhci/disk@g60a9800056716436636f393164695969.
```

```
33. c4t60A9800056716436636F393164695969d0 <DEFAULT cyl 9787 alt 2 hd 255
sec 63>
/scsi_vhci/disk@g60a9800056716436636F393164695969
```

3. Create the SAN file system that you will dump the current file system to.

This example uses the following command line to create a new file system.

```
newfs /dev/dsk/c4t60A9800056716436636F393164695969d0s0
```

4. Mount the file system that you will use when you copy the boot data.

The following example mounts the file system.

```
# mount /dev/dsk/c4t60A9800056716436636F393164695969d0s0 /mnt/bootlun
```

5. Create the required directory structure on the bootable LUN and copy the boot data. Enter:

```
ufsdump 0f - <source_boot_device> | cd /<mntpoint of bootable_ lun>;
ufsrestore rf -)
```

Note: If your configuration boots off of more than one device, you must create and configure a bootable LUN that matches each boot device on the host. If you are copying non-root boot partitions, see the Sun's document *Sun StorEdge SAN Foundation Software 4.4 Configuration Guide* for instructions.

The following example copies the information from c1d0s0.

```
# ufsdump 0f - /dev/rdisk/c1d0s0 | (cd /mnt/bootlun; ufsrestore rf -)
```

6. Edit the /etc/vfstab file to on the SAN file system. Change the swap, root, and file systems in the file to show the boot device instead of the local device. One way to do this is to edit the file using vi:

```
cd /mnt/bootlun/etc/
```

```
vi vfstab
```

The following example shows the vfstab entry for the bootable LUN.

#device	device	mount	FS	fsck	mount
#to mount	to fsck	point	type	pass	at boot
options					
#					
fd	-	/dev/fd fd	-	no	-
/proc	-	/proc proc	-	no	-
/dev/dsk/c4t60A9800056716436636F393164695969d0s1	-	-	-	-	swap
-	no	-	-	-	-
/dev/dsk/c4t60A9800056716436636F393164695969d0s0	/dev/rdisk/c4t60A9800056716436636F393164695969d0s0	/	ufs	1	no
-	-	-	-	-	-
/dev/dsk/c4t60A9800056716436636F393164695969d0s6	/dev/rdisk/c4t60A9800056716436636F393164695969d0s6	/globaldevices	ufs	2	yes
/devices	-	/devices	devfs	-	no

ctfs	-	/system/contract	ctfs	-	no	-
objfs	-	/system/object	objfs	-	no	-
swap	-	/tmp	tmpfs	-	yes	-

7. Modify the `/boot/solaris/bootenv.rc` file to reflect the new `bootpath` parameter. Use the boot device name noted above to replace the `bootpath`. To identify the boot slices, use "...a" at the end if your boot slice is slice 0, use "...b" for slice1, and so on.

Note: You can also see the device path by executing the `list` command:

```
ls -al /dev/rdisk/c4t60A9800056716436636F393164695969d0s2
```

The following example displays a modified file.

```
...
#
# Copyright 2005 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
#ident    "@(#)bootenv.rc 1.32    05/09/01 SMI"
#
#         bootenv.rc -- boot "environment variables"
#
setprop kbd-type US-English
setprop ata-dma-enabled 1
setprop atapi-cd-dma-enabled 0
setprop ttyb-rtts-dtr-off false
setprop ttyb-ignore-cd true
setprop ttya-rtts-dtr-off false
setprop ttya-ignore-cd true
setprop ttyb-mode 9600,8,n,1,-
setprop ttya-mode 9600,8,n,1,-
setprop lba-access-ok 1
setprop prealloc-chunk-size 0x2000
setprop bootpath /scsi_vhci/disk@g60a9800056716436636f393164695969:a
setprop console 'ttya'
```

8. Install GRUB on the new file systems.

```
cd /boot/grub; /sbin/installgrub stage1 stage2
/dev/rdisk/c3t500A09818350B481d32s0
```

Note: The GRUB configuration may indicate the wrong partition information for the bootable LUN. It may indicate that the boot slice is slice 1. If that is the case, you may need to change it to slice0 in the `menu.lst` file.

9. Update the GRUB bootloader:

```
bootadm update-archive -R /mnt/bootlun
```

10. Umount the file system on the bootable LUN.

```
umount /mnt/bootlun
```

11. Shut down the system.

sync;halt

12. Configure your HBA's bios to boot to the bootable LUN. Then configure your x86/x64 system to boot to the HBA. See the documentation for your server and HBA for more details.

x86/x64 without MPxIO systems: Copying data from locally booted server

If you have an x86/x64 system that is using MPxIO, you can copy data from a locally booted server.

Steps

1. Use the `mount` command to identify the current boot device.

The following example mounts the current boot device.

```
# mount / on /dev/dsk/c1d0s0
read/write/setuid/devices/intr/largefiles/logging/xattr/onerror=panic/dev=1980000
on Tue Jan 2 12:50:22 2007
```

2. Use the `format` command to choose the LUN. Write down the device path for the LUN so you can refer to it later. You need to configure the LUN for

- The root file system
- The file systems that will be copied to the bootable LUN. You may want to create slices and place the file systems on the slices.

In addition, you might want to set up the swap slice and slices for the metadata database replicas.

In this example, which uses a QLogic HBA, you have already run the `format` command and, based on its output, determined that you want number 33, which has the device path `/pci@0,0/pci10de,5d@e/pci1077,138@0,1/fp@0,0/disk@w500a09818350b481,20`.

```
c3t500A09818350B481d32 <DEFAULT cyl 9787 alt 2 hd 255 sec 63>
/pci@0,0/pci10de,5d@e/pci1077,138@0,1/fp@0,0/disk@w500a09818350b481,20
```

Note: Emulex HBAs have "fc" in the device name. QLogic HBAs have "fp" in the device name.

3. Create the SAN file system that you will dump the current file system to.

This example uses the following command line to create a new file system.

```
newfs /dev/dsk/c3t500A09818350B481d32s0
```

4. Mount the file system that you will use when you copy the boot data.

The following example mounts the file system.

```
mount /dev/dsk/c3t500A09818350B481d32s0 /mnt/boot lun
```

5. Create the required directory structure on the bootable LUN and copy the boot data. Enter:

```
ufsdump 0f - <source_boot_device> | (cd /<mntpoint of bootable_ lun>;  
ufsrestore rf -)
```

Note: If your configuration boots off of more than one device, you must create and configure a bootable LUN that matches each boot device on the host. If you are copying non-root boot partitions, see the Sun's document *Sun StorEdge SAN Foundation Software 4.4 Configuration Guide* for instructions.

The following example copies the information from c0t0d0s0.

```
ufsdump 0f - /dev/rdisk/c0t0d0s0 | (cd /mnt/bootlun; ufsrestore rf -)
```

6. Edit the /etc/vfstab file to on the SAN file system. Change the swap, root, and file systems in the file to show the boot device instead of the local device. One way to do this is to edit the file using vi:

```
cd /mnt/bootlun/etc/
```

```
vi vfstab
```

The following example shows the vfstab entry for the bootable LUN.

#device	device	mount	FS	fsck	mount
#to mount	to fsck	point	type	pass	at boot
options					
#					
fd -	/dev/fd fd	- no	-		
/proc -	/proc proc	- no	-		
/dev/dsk/c3t500A09818350B481d32s1	-	-	swap	-	no
-					
/dev/dsk/c3t500A09818350B481d320	/dev/rdisk/c3t500A09818350B481d32s0				
/	ufs 1	no	-		
/dev/dsk/c3t500A09818350B481d32s6	/dev/rdisk/c3t500A09818350B481d32s6				
/globaldevices	ufs 2	yes	-		
/devices	-	/devices	devfs	-	no
ctfs -	/system/contract	ctfs	-	-	no
objfs -	/system/object	objfs	-	no	-
swap -	/tmp	tmpfs	-	yes	-

7. Modify the /boot/solaris/bootenv.rc file to reflect the new bootpath parameter. Use the boot device name noted above to replace the bootpath. To identify the boot slices, use "...:a" at the end if your boot slice is slice 0, use "...:b" for slice1, and so on.

Note: You can also see the device path by executing the `list` command:

```
ls -al /dev/rdisk/c3t500A09818350B481d32s2
```

The following example sets the bootpath parameter.

```
...
setprop bootpath
/pci@0,0/pci10de,5d@e/pci1077,138@0,1/fp@0,0/disk@w500a09818350b481,20:a
...
```

8. Install GRUB on the new file systems.

```
cd /boot/grub; /sbin/installgrub stage1 stage2
/dev/rdisk/c3t500A09818350B481d32s2
```

Note: The GRUB configuration may indicate the wrong partition information for the bootable LUN. It may indicate that the boot slice is slice 1. If that is the case, you may need to change it to slice0 in the menu.lst file.

9. Update the GRUB bootloader:

```
bootadm update-archive -R /mnt/bootlun
```

10. Umount the file system on the bootable LUN.

```
umount /mnt/bootlun
```

11. Shut down the system.

```
sync;halt
```

12. Configure your HBA's bios to boot to the bootable LUN. Then configure your x86/x64 system to boot to the HBA. See the documentation for your server and HBA for more details.

Index

/etc/system

recommended values 77

A

ALUA

supported with MPxIO and FC 30

using with MPxIO 84

APM

example of installing 96

example of uninstalling 95

used with Veritas 89

available from Symantec 89

upgrading 90

used with Veritas 89

ASL

determining the ASL version 91

example of installing 96

example of uninstalling 95

upgrading 90

using VxVM 99

array type 98

available from Symantec 89

used with Veritas 89

B

basic_config command

options with LPFC drivers 77

options with MPxIO 86

options with native drivers and DMP 80

installed by Host Utilities 24

native drivers and Veritas options 81

options with native drivers and Veritas 81

Veritas DMP environments 74

bindings, See persistent bindings

brocade_info

collecting switch information 144

installed by Host Utilities 24

C

cfgadm -al command 114

CHAP

configuring 71

secret value 71

cisco_info

collecting switch information 144

installed by Host Utilities 24

configurations

supported for Host Utilities 27

finding more information 33

controller_info

collecting storage system information 143

installed by Host Utilities 24

create_binding program

binding WWPNS 147

creating persistent bindings 147

executing 149

installed by Host Utilities 24

D

device names

generated using LUNs serial numbers 133

diagnostic utilities

brocade_info 144

cisco_info 144

controller_info 143

executing 142

filer_info 143

mcddata_info 144

qllogic_info 144

solaris_info 142

documentation

finding more information 33

drivers

getting hbacmd 43

getting software 41

installing SANsurfer CLI 53

dynamic discovery

iSCSI SendTargets command 68

E

Emulex

changing HBA to SD mode 160

Emulex (*continued*)

- changing HBA to SFS mode [155](#)
- downloading software and firmware [42](#)
- getting software [41](#)
- LPFC drivers 6.21g [42](#)
- LPFC drivers prior to 6.21g [42](#)
- using Sun native drivers [50](#)

error messages

- ASL [103](#)

example

- basic_config command with LPFC drivers [78](#)
- basic_config command with MPxIO [87](#)
- basic_config command with native drivers and Veritas [82](#)
- installing ASL, APM [96](#)
- uninstalling ASL, APM [95](#)

F

fast recovery

- Veritas feature [102](#)

FC protocol [26, 30](#)

- ALUA and MPxIO [30](#)

FCA utilities

- EMLXemlxu [50](#)

filer_info

- collecting storage system information [143](#)
- installed by Host Utilities [24](#)

finding more information [33](#)

format utility

- labeling LUNs [115](#)

H

HBA

- displaying information with sanlun [123](#)

hbacmd

- creating persistent bindings [110](#)

host

- collecting information using solaris_info [142](#)

Host Utilities

- contents [24](#)
- downloading software packages [58](#)
- environments [21](#)
- finding more information [33](#)
- general componets [24](#)
- installing [59](#)
- key setup steps [57](#)
- MPxIO environment [21](#)
- MPxIO stack overview [32](#)

Host Utilities (*continued*)

- software packages [58](#)
- support for non-English versions of Solaris [31](#)
- uninstalling Attach Kit 2.0 [65](#)
- uninstalling versions 5.x, 4.x, 3.x [64](#)
- upgrading [63](#)
- Veritas environment [21](#)
- Veritas stack overview [31](#)

I

igroup create command [108](#)

information

- finding more [33](#)

installation

- iSCSI configuration [67](#)
- downloading software packages [58](#)
- key setup steps [57](#)

iSCSI protocol [26, 29, 67, 68, 69, 70, 71, 113, 139](#)

- configuration [67](#)
- storage system IP address [68](#)
- configuring bidirectional CHAP [71](#)
- configuring discovery [68](#)
- disabling MPxIO [70](#)
- discovering LUNs [113](#)
- MPxIO [29](#)
- node names [67](#)
- recording node names [68](#)
- troubleshooting [139](#)
- Veritas DMP [29](#)
- working with Veritas [69](#)

ISNS discovery

- iSCSI [68](#)

L

labeling LUNs [115](#)

languages

- support for non-English versions of Solaris [31](#)

LPFC drivers

- example of basic_config command [78](#)
- getting hbacmd [43](#)
- adding target ID entries [111](#)
- discovering LUNs with 6.21g drivers [112](#)
- discovering LUNs with drivers prior to 6.21g [112](#)
- getting Emulex software [41](#)
- SAN booting [160](#)

lun create command [108](#)

lun map command [108](#)

LUNs

- configuration overview 105
- creating 108
- creating Veritas boot LUN 153
- discovering when using iSCSI 113
- discovering with 6.21g drivers 112
- discovering with drivers prior to 6.21g 112
- discovering with native drivesr 115
- displaying with sanlun 119
- labeling 115

M

- man pages
 - installed by Host Utilities 24
- mcdata_info
 - collecting switch information 144
 - installed by Host Utilities 24
- mpathadm
 - verifying multipathing on MPxIO 133
- MPxIO
 - example of basic_config command 87
 - ALUA 21
 - checking multipathing 133
 - environment for Host Utilities 21
 - fast recovery with Veritas 102
 - iSCSI protocol 29
 - labeling LUNs 115
 - multipathing 21
 - protocols 21
 - removing symmetric-option 84
 - sd.conf variables 85
 - ssd.conf variables 85
 - stack overview 32
 - troubleshooting 145
 - using ALUA 84
- mpxio_set script
 - installed by Host Utilities 24
 - preparing for ALUA 84
 - removing symmetric-option 84
- multipathing
 - verifying on MPxIO 133

N

- node names
 - iSCSI 67
 - recording 68
- non-English versions of Solaris supported 31

O

- outpassword
 - CHAP secret value 71

P

- paths
 - displaying with sanlun 121
- persistent bindings
 - creating with hbacmd 110
- polling interval
 - recommended values 102
- problems
 - checking troubleshooting 128
- publications
 - finding more information 33

Q

- QLogic
 - qlc 52
 - creating FCode compatibility 158
 - getting HBA software 41
- qllogic_info
 - collecting switch information 144
 - installed by Host Utilities 24

R

- Release Notes
 - checking 127
- requirements
 - finding more information 33
- restore policy
 - recommended value 102

S

- SAN booting
 - changing Emulex HBA to SD mode 160
 - changing Emulex HBA to SFS mode 155
 - copying data on x86/64 without MPxIO 185
 - direct installations with native drivers 162
 - FCodes compatibility with QLogic HBA 158
 - MPxIO environments 177
- san_version command
 - installed by Host Utilities 24

- sanlun utility
 - displaying LUNs [119](#)
 - displaying multipathing for Veritas [101](#)
 - displaying paths [121](#)
 - installed by Host Utilities [24](#)
 - verifying multipathing on MPxIO [133](#)
- sd.conf
 - adding target ID entries [111](#)
 - recommended values for MPxIO [85](#)
- software packages
 - downloading Host Utilities software [58](#)
 - installing Host Utilities [59](#)
 - SPARC processor [58](#)
 - x86/64 processor [58](#)
- Solaris
 - Host Utilities [21](#)
 - Host Utilities environments [21](#)
 - labeling LUNs [115](#)
 - support for non-English language versions [31](#)
- solaris_efi
 - LUN type [108](#)
- solaris_info
 - collecting host information [142](#)
 - installed by Host Utilities [24](#)
- SPARC processor
 - software package [58](#)
- ssd_config.pl script
 - installed by Host Utilities [24](#)
- ssd.conf file
 - recommended values with Veritas DMP [80](#)
- static discovery
 - iSCSI [68](#)
- storage system
 - outpassword and CHAP [71](#)
- storage systems
 - displaying path information with sanlun [121](#)
 - collecting information using controller_info [143](#)
- Sun HBAs
 - working with [41](#)
- Sun native drivers
 - example of basic_config command with MPxIO [87](#)
 - example of basic_config command with Veritas [82](#)
 - basic_config command options with Veritas [81](#)
 - direct SAN boot installations [162](#)
 - FCA utilities [50](#)
 - getting Emulex HBAs [41](#)
- SVM
 - managing LUNs [117](#)
 - MPxIO [29](#)

- switch
 - collecting information using brocade_info [144](#)
 - collecting information using cisco_info [144](#)
 - collecting information using mcdata_info [144](#)
 - collecting information using qllogic_info [144](#)
- Symantec
 - provides ASL, APM [89](#)
- symmetric-option
 - removing for ALUA [84](#)

T

- troubleshooting
 - checking Release Notes [127](#)
 - finding more information [33](#)

V

- Veritas
 - fast recovery [102](#)
 - restore daemon settings [102](#)
 - volume manager [29](#)
 - APM [89](#)
 - ASL [89](#)
 - basic_config command [74](#)
 - basic_config command options with native drivers [81](#)
 - configuring iSCSI [69](#)
 - drivers [21](#)
 - environment for Host Utilities [21](#)
 - getting driver software [41](#)
 - iSCSI protocol [29](#)
 - labeling LUNs [115](#)
 - multipathing [21](#)
 - preparing for Host Utilities [74](#)
 - protocols [21](#)
 - setup issues [21](#)
 - stack overview [31](#)
- volume management
 - managing LUNs [117](#)
- VxVM
 - displaying LUN paths [99](#)
 - Veritas volume manager [29](#)
 - managing LUNs [117](#)

W

- WWPNs
 - binding to target ID [147](#)

X

- x86/64
 - copying data for SAN boot LUN without MPxIO [185](#)
- x86/64 processor
 - software package [58](#)

Z

- ZFS
 - managing LUNs [117](#)

