

SAN Host Utilities 4.2 for Fibre Channel on Linux®

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: 215-04337_A0
December 2008

Copyright and trademark information

Copyright information

Copyright © 1994-2009 NetApp, Inc. All rights reserved. Printed in the U.S.A.

The third-party copyright statements and licenses that apply to open source software components distributed with the product are contained in the section “Third-Party Copyrights and Licenses.”

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

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

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

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

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

Trademark information

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

Design; SharedStorage; Simplicore; Simulate ONTAP; Smart SAN; SnapCache; SnapDirector; SnapFilter; SnapMigrator; SnapSuite; SohoFiler; SpinMirror; SpinRestore; SpinShot; SpinStor; vFiler; 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. NetApp, Inc. NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	Overview	1
	Understanding the Host Utilities components and requirements.	3
	Understanding host and storage system communication	6
	The setup procedure	8
	Where to find the information you need	11
Chapter 2	Installing HBAs and Software on the Host	15
	Installing the FC Linux Host Utilities software	16
	Removing the previous external HBA drivers	18
	Verifying or installing the required multipathing packages	20
	Installing HBAs.	21
	Downloading and installing the external HBA driver and applications	22
	Setting the required HBA and driver parameters.	24
	Bundled HBA drivers and applications.	27
	Recording the WWPNs of the HBAs.	34
Chapter 3	Configuring the Storage System for FC	37
Chapter 4	Migrating a Storage System to single_image mode	39
	Understanding cfmode	40
	Migrating to single_image mode	42
Chapter 5	Configuring FC Cabling and Switches.	47
	Switched fabric configurations	48
	Direct-connected configurations	50

Chapter 6	Accessing LUNs from a Red Hat Enterprise Linux Host.	51
	Loading the external HBA driver on the host	53
	Loading the bundled HBA driver on the host	54
	Configuring dm-multipath	55
	Starting the multipath service.	61
	Accessing LUNs with dm-multipath support.	62
	Verifying the HBA setting and dm-multipath configuration.	63
	Identifying and viewing a list of LUNs	66
	Creating and mounting a file system on a multipath device	68
	Tuning the dm-multipath configuration	70
	Stopping dm-multipath	71
	Accessing LUNs with Veritas DMP support.	72
	Accessing LUNs without dm-multipath support	73
	Identifying and viewing LUN information.	75
	Creating a partition	76
	Creating a file system	78
	Labeling and mounting the file system.	79
	Unloading the external HBA driver and shutting down	81
	Unloading the bundled HBA driver and shutting down	82
 Chapter 7	 Accessing LUNs from a SUSE Linux Enterprise Server Host	 83
	Loading the external HBA driver on the host	84
	Loading the bundled HBA driver on the host	85
	Configuring dm-multipath	86
	Starting multipath service.	91
	Accessing LUNs with dm-multipath support.	92
	Verifying the HBA setting and dm-multipath configuration.	93
	Identifying and viewing a list of LUNs	96
	Creating and mounting a file system on a LUN	98
	Tuning the dm-multipath configuration	100
	Stopping dm-multipath	101
	Accessing LUNs with Veritas DMP support.	102
	Accessing LUNs without dm-multipath support	103
	Identifying and viewing the LUN information.	105
	Creating a partition	107

	Creating a file system	108
	Labeling and mounting the file system.	109
	Unloading the external HBA driver and shutting down	111
Chapter 8	Configuring SAN Boot on a Red Hat Enterprise Linux Host	113
	Configuring SAN boot on the native host	114
	Configuring SAN boot for Veritas Storage Foundation	121
Chapter 9	Configuring SAN Boot on a SUSE Linux Enterprise Server Host.	125
	Configuring SAN boot on the native host	126
	Configuring SAN boot for Veritas Storage Foundation	131
Chapter 10	Veritas Array Support Library and Array Policy Module.	133
	Overview of the Array Support Library and Array Policy Module	134
	Installing the ASL and APM	135
	Removing the ASL and APM	137
	Obtaining multipathing information	139
	Understanding ASL error message types.	143
Appendix A	Troubleshooting	145
Appendix B	Troubleshooting—Veritas Storage Foundation	151
Appendix C	Third-Party Copyrights and Licenses	153

Preface

About this guide

This guide describes how to install the FC Linux® Host Utilities 4.2 for Native OS and Veritas® on a Linux host and how to set up the host to access the storage object described by a logical unit number (LUN). It also explains how to configure the storage system to work with the host.

For more information about specific Linux versions supported, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Audience

This guide is for administrators of Red Hat® Enterprise Linux, SUSE® Linux Enterprise Server, Oracle® Enterprise Linux or Veritas Storage Foundation™ hosts, and NetApp storage systems.

Terminology

This guide uses the following terms:

- ◆ NetApp storage products (storage systems, FAS appliances, and NearStore® systems) are all *storage systems*.
- ◆ The term *type* means pressing one or more keys on the keyboard.

Command conventions

You can enter storage system commands on the system console or from any client that can obtain access to the storage system by using a Telnet session. In examples that illustrate commands executed on a Linux workstation, the command syntax and output might differ, depending on your version of Linux.

Formatting conventions

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

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

Keyboard conventions

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

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

Special messages

This guide contains special messages that are described as follows:

Note

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

Attention

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

About this chapter

This chapter contains important information you should read before beginning the setup process.

About Linux versions

FC Linux Host Utilities supports the Red Hat Enterprise Linux, SUSE Linux Enterprise Server, Oracle Enterprise Linux operating systems, and Veritas Storage Foundation with Red Hat Enterprise Linux and SUSE Linux Enterprise Server. For information about the supported operating system version and Veritas Storage Foundation version, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Where the setup procedure differs based on the Linux version, the relevant information is called out in this guide. If not otherwise marked, the information in this guide applies to all supported versions of Linux software.

Note

The setup procedures are the same for Red Hat Enterprise Linux and Oracle Enterprise Linux.

About host bus adapters

FC Linux Host Utilities supports both QLogic® and Emulex® brand host bus adapters (HBAs).

With this version of Host Utilities, NetApp supports the bundled HBA drivers along with the external HBA drivers while using dm-multipath on Linux FC hosts.

The term *external HBA driver* is used to specify the drivers that you must download separately from the HBA vendor's Web site.

The term *bundled HBA driver* is used to specify the drivers that come bundled along with the operating system.

The setup procedures for the bundled drivers are mentioned separately in the guide.

Where the setup procedure differs based on HBA model, the relevant information is called out in this guide. If not otherwise marked, the information in this guide applies to all supported versions of HBA models.

Note

Do not use different HBA brands on a single Linux host.

About SAN boot

SAN boot uses a SAN-attached disk, such as a LUN configured on storage controller, as a boot device for a host. For information about the supported operating system version and Veritas Storage Foundation version, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Understanding the Host Utilities components and requirements](#)” on page 3
- ◆ “[Understanding host and storage system communication](#)” on page 6
- ◆ “[The setup procedure](#)” on page 8
- ◆ “[Where to find the information you need](#)” on page 11

Understanding the Host Utilities components and requirements

What FC Linux Host Utilities is

FC Linux Host Utilities 4.2 enables a Linux host to connect to your storage systems. FC Linux Host Utilities consists of the following components:

- ◆ Utilities, scripts, and support files
- ◆ Documentation, including this *Setup Guide*, *Release Notes*, and *Quick command reference* guides

Host Utilities components

FC Linux Host Utilities 4.2 consists of the following software components:

- ◆ The `sanlun` utility that displays information about LUNs on the storage system that are available to this Linux host
- ◆ The `install` and `uninstall` command scripts to add and remove the FC Linux Host Utilities components
- ◆ The following diagnostic scripts for collecting troubleshooting information:
 - ❖ `linux_info`: collects configuration information about your Linux host including Veritas stack information
 - ❖ `filer_info`: collects information about the storage system
 - ❖ `qlogic_info`: collects information about QLogic® switches installed in the network
 - ❖ `brocade_info`: collects information about Brocade® switches installed in the network
 - ❖ `cisco_info`: collects information about Cisco® switches installed in the network
 - ❖ `mcddata_info`: collects information about McData® switches installed in the network
 - ❖ `san_version`: displays the version of the Host Utilities
 - ❖ `SHsupport.pm`: support library for `*info` scripts
 - ❖ `Telnet.pm`: support library for `*info` scripts
 - ❖ `FTP.pm`: support library for `*info` scripts
- ◆ The manual (`man`) pages for `sanlun` and the diagnostic scripts
- ◆ This *Setup Guide*
- ◆ The *Release Notes*
- ◆ The Quick command reference guides

Native multipathing

Linux device mapper multipathing (dm-multipath) is supported on storage systems with Linux hosts running Red Hat Enterprise Linux 4 Update 2 and later, SUSE Linux Enterprise Server 9 SP2 and later, and Oracle Enterprise Linux 4 and later. This support enables you to configure multiple network paths between the Linux host and storage system. If one path fails, FC traffic continues on the remaining paths. The dm-multipath support is required if the host has multiple paths to a LUN.

The required kernel modules for dm-multipath are included in Red Hat Enterprise Linux 4 Update 2 and later, SUSE Linux Enterprise Server 9 SP2 and later, and Oracle Enterprise Linux 4 and later binary kernel packages. You might have to install additional user space packages as part of the setup process.

The specific steps for enabling dm-multipath are included in this guide.

Veritas DMP

Veritas Volume Manager (VxVM) uses Dynamic Multipathing (DMP) to provide multipathing. For more information about Veritas Volume Manager and the procedure for accessing LUNs, see the *Veritas Volume Manager Administrator's Guide* for Linux that is shipped along with the software.

ALUA support

The FC Linux Host Utilities adds support for Asymmetric Logical Unit Access (ALUA). To see which versions of Data ONTAP® and Host Utilities supports ALUA, see the *Compatibility and Configuration Guide* for NetApp's FC and iSCSI Products at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

NetApp introduced ALUA support with Data ONTAP 7.2 and Single Image cfmode.

Note

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

ALUA defines a standard set of SCSI commands for discovering and managing multiple paths to LUNs on Fibre Channel and iSCSI SANs. You still need native multipathing on the host for ALUA support.

Note

ALUA is not supported with Veritas Storage Foundation.

Data ONTAP requirements

See the *Compatibility and Configuration Guide for NetApp's FC and iSCSI Products* for supported versions of Data ONTAP. If you do not have a supported version of Data ONTAP, upgrade to one of the supported versions.

When you upgrade to a newer version of Data ONTAP, the cfmode setting remains the same. If you want to change the cfmode, for example to enable communication with the 4-Gb target HBA, then you must manually change the cfmode.

See Chapter 4, “[Migrating a Storage System to single_image mode](#),” on page 39 for more information.

Understanding host and storage system communication

About host and storage system communication

After installing and configuring your Fibre Channel HBA, the Linux host uses the FC for SCSI to access data on storage systems that run supported versions of Data ONTAP software.

The storage system is an FC target device. Each host HBA is an initiator that uses FC to access LUNs on a storage system to store and retrieve data.

On a host, a WWPN identifies each port on an HBA. The host WWPNs are used as identifiers when creating igroups on a storage system. An initiator group permits the host access to specific LUNs on a storage system.

Supported configurations

The following configurations are supported:

- ◆ Fabric-attached SAN

Two variations of fabric-attached SANs are supported:

- ❖ 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 a high-availability storage system (active/active configuration). A fabric-attached single-path host has one HBA.

- ❖ A dual-host FC connection from the HBA to the storage system through dual switches

The redundant configuration avoids the single point of failure of a single-switch configuration.

- ◆ Direct-attached

In this context, “direct-attached” means a single host with a direct FC connection from the host HBA to stand-alone or high-availability (active/active) storage system configurations.

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.

For more information (including diagrams) about the supported Fibre Channel topologies, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

The setup procedure

Before you begin

Use the appropriate interoperability matrices to verify that you have a supported combination of the following:

- ◆ Linux operating system version
- ◆ HBA model
- ◆ HBA firmware and driver version
- ◆ Fibre Channel switch model and firmware version
- ◆ Storage system and Data ONTAP version
- ◆ Veritas Storage Foundation

You can find the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Setup procedure

The procedure for configuring the Fibre Channel (FC) protocol on a host and storage system follows the same basic sequence for all hosts.

Step	Action
1	Install and configure HBAs and software on the host, including the following: ◆ Installing the HBAs if not already installed ◆ Removing any existing external HBA drivers ◆ Downloading and installing the supported external HBA driver ◆ Installing the FC Linux Host Utilities software ◆ Setting required HBA parameters ◆ Installing dm-multipath packages to use the Native multipathing solution ◆ Recording the worldwide port names (WWPNs) of the HBAs ◆ Installing Veritas Storage Foundation, if you want to use the Veritas DMP solution

Step	Action
	Note _____ For information about installing Veritas Storage Foundation, see the <i>Veritas Storage Foundation Installation Guide</i> for Linux that is shipped along with the software. ❖ Install Array Support Library (ASL) and Array Policy Module (APM) See Chapter 10, “Veritas Array Support Library and Array Policy Module,” on page 133. See Chapter 2, “Installing HBAs and Software on the Host,” on page 15.
2	Configure the storage system, including the following: ◆ Licensing and starting the FC service ◆ Creating LUNs, creating an initiator group (igroup) that contains the host’s FC WWPN, and mapping the LUNs to that igroup Note _____ The LUN 0 must be mapped to the host. ◆ Optional: Enable ALUA, if you have not already done so. See Chapter 3, “Configuring the Storage System for FC,” on page 37.
3	Optional: Migrate an active/active storage system configuration to single_image cfmode. See Chapter 4, “Migrating a Storage System to single_image mode,” on page 39.
4	Configure the Fibre Channel switches and cabling, including the following: ◆ Verifying the firmware revision of the switches ◆ Cabling the storage system to the switches ◆ Creating zones on the switches ◆ Cabling the HBAs to the switches or directly to the storage system after the LUN 0 is mapped to the host See Chapter 5, “Configuring FC Cabling and Switches,” on page 47.

Step	Action
5	Discover and access the LUNs from the host, including the following: ◆ Loading the external HBA driver modules ◆ Configuring multipathing using dm-multipath or Veritas DMP See Chapter 6, “Accessing LUNs from a Red Hat Enterprise Linux Host,” on page 51, or Chapter 7, “Accessing LUNs from a SUSE Linux Enterprise Server Host,” on page 83.

Where to find the information you need

About storage area network (SAN) documentation

This setup guide does not include detailed instructions for setting up and configuring your storage system. However, the following section describes some important SAN-related documents that you should consult to ensure your host and storage system are properly configured.

Where to go for host and storage system information

The following table lists documents with the most current information about host and storage system requirements, supported configurations, and cabling instructions.

If you need information about...	See...
Known issues, system requirements, and last minute updates	The <i>Release Notes</i> included with this Host Utilities, available at http://now.netapp.com/NOW/knowledge/docs/san/ .
The latest supported configurations	◆ The NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml. ◆ The <i>System Configuration Guide</i> at https://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/.
Configuring the storage system	◆ The <i>Best Practices for Reliability</i> at http://now.netapp.com/NOW/knowledge/docs/docs.cgi. ◆ The <i>Data ONTAP Software Setup Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

If you need information about...	See...
Supported Fibre Channel SAN topologies	The NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability , or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml .
Installing and configuring the HBA in your host	Respective HBA vendor documentation.
Installing and configuring a supported version of SnapDrive for UNIX	The <i>Installation and Administration Guide</i> for that version of SnapDrive for UNIX.
Managing SAN storage on the storage system	◆ The <i>Command Reference</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ The <i>Data ONTAP Block Access Management Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ The FilerView online Help ◆ The <i>Best Practices for Reliability: Ongoing Maintenance</i> at http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm.
Installing Veritas Storage Foundation	The <i>Veritas Storage Foundation™ Installation Guide</i> for Linux that is shipped along with the software.
Release information about the products in the Veritas Storage Foundation Linux product line	The <i>Veritas Storage Foundation™ Release Notes</i> for Linux that is shipped along with the software.

If you need information about...	See...
Veritas Volume Manager	◆ The <i>Veritas™ Volume Manager Administrator's Guide</i> for Linux that is shipped along with the software. ◆ The <i>Veritas Volume Manager Troubleshooting Guide</i> for Linux that is shipped along with the software.
Veritas File System	◆ The <i>Veritas File System Administrator's Guide</i> for Linux that is shipped along with the software.
Veritas Cluster Server	◆ The <i>Veritas™ Cluster Server Installation Guide</i> for Linux that is shipped along with the software. ◆ The <i>Veritas™ Cluster Server User's Guide</i> for Linux that is shipped along with the software. ◆ The <i>Veritas™ Cluster Server Release Notes</i> for Linux that is shipped along with the software.
Veritas Storage Foundation for Oracle RAC	◆ The <i>Veritas Storage Foundation™ for Oracle RAC Installation and Configuration Guide</i> for Linux that is shipped along with the software. ◆ The <i>Veritas Storage Foundation™ for Oracle RAC Release Notes</i> for Linux that is shipped along with the software.
Veritas Storage Foundation Cluster File System	◆ The <i>Veritas Storage Foundation™ Cluster File System Installation Guide</i> for Linux that is shipped along with the software. ◆ The <i>Veritas Storage Foundation™ Cluster File System Administrator's Guide</i> for Linux that is shipped along with the software. ◆ The <i>Veritas Storage Foundation™ Cluster File System Release Notes</i> for Linux that is shipped along with the software.

If you need information about...	See...
Veritas Storage Foundation Cluster File System for Oracle RAC	◆ The <i>Veritas Storage Foundation™ Installation Guide</i> for Linux that is shipped along with the software. ◆ The <i>Veritas Storage Foundation™ Cluster File System Administrator's Guide</i> for Linux that is shipped along with the software. ◆ The <i>Veritas Storage Foundation™ Cluster File System Release Notes</i> for Linux that is shipped along with the software.
Enabling ALUA	<i>Data ONTAP Block Access Management Guide</i> for FC at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. In particular, see the section Enabling ALUA for a Fibre Channel igroup.

Installing HBAs and Software on the Host

Configuration overview

This chapter explains how to install the FC Linux Host Utilities, HBAs, and drivers, and how to complete the initial configuration of the Linux host. Complete the final configuration of the host after you have configured the storage system.

If you want to use the Linux dm-multipath support to create a highly-available connection between the Linux host and the storage system, you also need the device-mapper (Red Hat Enterprise Linux and SUSE Enterprise Linux Server) and multipath-tools (SUSE Enterprise Linux Server) or device-mapper-multipath (Red Hat Enterprise Linux) packages.

If you want to use Veritas DMP support for highly-available connection between the Linux host and storage system, you also need to install Veritas Storage Foundation ASL and APM packages.

Note

Do not connect the Fibre Channel cables to the HBAs on the Linux host until you have created LUN 0 and mapped it to the host.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Installing the FC Linux Host Utilities software”](#) on page 16
- ◆ [“Removing the previous external HBA drivers”](#) on page 18
- ◆ [“Verifying or installing the required multipathing packages”](#) on page 20
- ◆ [“Installing HBAs”](#) on page 21
- ◆ [“Downloading and installing the external HBA driver and applications”](#) on page 22
- ◆ [“Setting the required HBA and driver parameters”](#) on page 24
- ◆ [“Bundled HBA drivers and applications”](#) on page 27
- ◆ [“Recording the WWPNs of the HBAs”](#) on page 34

Requirement for root privileges

Before you configure the initiator, ensure you are logged in with root privileges.

Installing the FC Linux Host Utilities software

Verifying the correct version

Verify that you have the correct version of FC Linux Host Utilities for your version of Linux. For information about the correct version, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Copying or downloading the Host Utilities software

Copy the FC Linux Host Utilities file (netapp_linux_host_utilities_4_2.tar.gz) from the CD or download it from the NOW (NetApp on the Web)[®] site at <http://now.netapp.com/NOW/cgi-bin/software> to a working directory on your Linux host.

Note

In SUSE Linux Enterprise Server 9 SP3 x86_64, there is no 32-bit compatibility library for libdevmapper available. The 32-bit compatibility package for device-mapper-multipath can be downloaded from the Novell[®] Web site at <http://support.novell.com/techcenter/psdb/8ae2a27b3b63ac95fd776298026596eb.html>. This package needs to be installed, to ensure that sanlun works on SUSE Linux Enterprise Server 9 SP3 x86_64 bit.

Installing the Host Utilities software

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

Step	Action
1	Remove any previous version of FC Linux Host Utilities. Go to the directory where the previous version is installed (default is /opt/netapp/santools) and enter the following command: ./uninstall
2	Go to the working directory to which you copied the Host Utilities file.
3	To uncompress the file, enter the following command: gunzip netapp_linux_host_utilities_4_2.tar.gz

Step	Action
4	To extract the files, enter the following command: tar -xvf netapp_linux_host_utilities_4_2.tar
5	Change to the netapp_linux_host_utilities_4_2 directory. By default, this is a subdirectory of the working directory in which you extracted the Host Utilities files in the previous step.
6	Enter the following command: ./install

For detailed information about running the diagnostic scripts, see the man pages in the /opt/netapp/man/man1 directory.

Removing the previous external HBA drivers

Removing previous external QLogic HBA drivers

You need to remove any previous external HBA driver modules. To remove the QLogic driver modules, complete the following steps.

Step	Action
1	Log in to the Linux host as root.
2	Stop the qlremote process if it is running by entering the following command: pskill qlremote
3	Change to the directory where the driver installer files are located.
4	Uninstall the previous QLogic driver on the host. Ensure that the driver is not in use (that is, usage count of qla2xxx is 1, because it is referenced by qla2300/qla2400/qla2500 alone). ◆ For a SANsurfer package that is installed on the host, run the following command: ./qlinstall --uninstall ◆ For a freshly installed host with the bundled drivers loaded, run the following command: modprobe -r qla2300/qla2400/qla2500

Removing previous external Emulex HBA drivers

You need to remove any previous external HBA driver modules. To remove the Emulex driver modules, complete the following steps.

Step	Action
1	Log in to the Linux host as root.
2	Go to the directory where the driver installer files are located.

Step	Action
3	Uninstall the previous Emulex driver on the host. Ensure that the driver is not in use (that is, usage count of lpfc is 0). ◆ For a freshly installed host with the bundled drivers loaded, run the following command: modprobe -r lpfc ◆ Run the following command: ./lpfc-install --uninstall

Verifying or installing the required multipathing packages

Multipathing packages

If you want to use the Linux dm-multipath support to create a highly available connection between the Linux host and storage system, you need to verify that the correct multipathing packages are installed on your Linux host. To verify this, complete the following steps.

Step	Action
1	If you plan to use dm-multipath support, enter one of the following commands: rpm -q device-mapper (both SUSE Linux Enterprise Server and Red Hat Enterprise Linux) rpm -q multipath-tools (SUSE Linux Enterprise Server hosts only) rpm -q device-mapper-multipath (Red Hat Enterprise Linux hosts only) Result: The rpm -q commands return the name and version of the installed packages.
2	If you do not have the required packages, install them by using the appropriate method for your version of Linux.

If you want to use the Veritas DMP support to create a highly available connection between the Linux host and storage system, ensure that the Storage Foundation Software and ASL/APM are installed properly. For more information, see “[Accessing LUNs with Veritas DMP support](#)” on page 72 and “[Veritas Array Support Library and Array Policy Module](#)” on page 133.

Installing HBAs

Verify HBA model is supported

Verify that your HBA model is supported. For information about the supported HBA models, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Do not mix HBA models

If you install more than one HBA, be sure that all HBAs are the same brand and bus type (PCI-X or PCI Express). Mixing HBA brands on one host can lead to driver conflicts.

Hardware installation

Follow the HBA and host hardware vendors' instructions for installing the HBAs on the Linux host, if the HBAs are not already installed.

Downloading and installing the external HBA driver and applications

Verify required driver version

The HBA requires a specific driver version that is customized for use with a storage system. Be sure to verify the required driver version by using the support matrix before downloading. For details about the supported driver versions, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

You also need to install the HBA vendor's appropriate management application.

Download from vendor's Web site

You must download the external HBA driver software package, applications package, and the rescan script from the HBA vendor's Web site.

QLogic: For QLogic HBAs, go to http://support.qlogic.com/support/drivers_software.asp. Under OEM Models, select NetApp. Find the driver version listed in the support matrix and download it.

Emulex: For Emulex HBAs, go to <http://www.emulex.com/ts/index.html>. Under Storage and System Supplier Qualified and Supported HBAs, select NetApp. Find the driver version listed in the support matrix and download it. Also download the Emulex applications package from the same location.

Install the driver software

Follow the instructions on the vendor's Web site or in the README included with the driver for installing the driver software on your Linux host.

After installing the external HBA driver package, reboot the system.

QLogic HBA: Run the `qlinstall` script with no options to install the driver package.

Emulex HBA: Run the `lpfc-install` script with no options to install the driver package.

Install HBA management software

You should install the HBA vendor's command-line interface (CLI) and graphical user interface (GUI) management packages. These packages contain software components used to display HBA information (also used by the `sanlun` utility).

QLogic: Download and install the following from the QLogic Web site:

- ❖ SANsurfer FC HBA Manager (for `sanlun fcp show adapter`)
- ❖ SANsurfer CLI

Emulex: Download and install the following from the Emulex Web site:

- ❖ Applications kit: Ensure that the application kit version matches the bundled Emulex driver version.

Setting the required HBA and driver parameters

For QLogic HBAs

You need to set both HBA and driver parameters to ensure that the multipathing and storage system failover work correctly.

QLogic HBA parameters: To set the HBA parameters on the Linux host by using the QLogic **scli** command, complete the following steps.

Step	Action
1	Run the scli command: scli
2	From the main menu, select the Configure HBA Settings option.
3	Select the first HBA port listed.
4	Configure the QLogic HBA port topology to Point to Point Only while connecting to a storage system operating in SSI, Standby or Partner CF modes, or to Loop Only if the storage system is operating in Dual Fabric or Mixed CF modes. This can be done by selecting ConnectionOptions and setting it to Point to Point Only or Loop Only, depending on the topology that you are using.
5	Configure the QLogic port speed ideally to the highest speed possible, depending on its ceiling, and that of the switch or target port it is connected to. To do this select the Data Rate option and set it to the specified speed.
6	Select the option for Port Down Retry Count and set the value to 30.
7	Select the option for Link Down Timeout and set the value to 20.
8	Select Commit Changes for this HBA so that the changes are applied to the HBA ports.
9	Repeat the settings for each HBA port listed.
10	Return to the main menu and select Quit.

QLogic driver parameters: To set the driver parameters, complete the following steps on the Linux host.

Note

You need to edit either `/etc/modprobe.conf` (Red Hat Enterprise Linux) or `/etc/modprobe.conf.local` (SUSE Linux Enterprise Server), depending on where the QLogic installer put its options.

Step	Action
1	Unload the external QLogic HBA driver by entering the appropriate command for your HBA model: <code>./qlinstall -ul qla2500</code> <code>./qlinstall -ul qla2400</code> <code>./qlinstall -ul qla2300</code>
2	Append the <code>options qla2xxx</code> statement with settings in the <code>modprobe.conf</code> file on a single line; only one <code>options qla2xxx</code> statement is allowed. <code>ql2xsuspendcount=6 MaxRetriesPerPath=10</code> <code>MaxRetriesPerIo=10 ql2xfailover=0</code>
3	Save the file and exit.
4	Reload the external QLogic HBA driver using the appropriate command for your HBA model: <code>./qlinstall -l qla2500</code> <code>./qlinstall -l qla2400</code> <code>./qlinstall -l qla2300</code>
5	Enter the appropriate command for your HBA model to rebuild the initial ramdisk (initrd) image with the new parameter values: <code>./qlinstall -br -in qla2500</code> <code>./qlinstall -br -in qla2400</code> <code>./qlinstall -br -in qla2300</code>
6	Reboot the Linux host by using the updated image.

For Emulex HBAs

No changes are required.

Note

Disable subfs and hwscan for SUSE Linux Enterprise Server 9 by setting `HOTPLUG_USE_HWSCAN=no` and `HOTPLUG_USE_SUBFS=no` in `/etc/sysconfig/hotplug`.

Bundled HBA drivers and applications

Installing HBA management software

You should download and install the HBA vendor’s command-line interface (CLI) and graphical user interface (GUI) management packages from the respective HBA vendor Web site. These packages contains software components used to display HBA information (also used by the sanlun utility). You should also download the rescan script from the HBA vendor Web site.

QLogic: Download and install the following from the QLogic Web site:

- ❖ SANsurfer FC HBA Manager (for `sanlun fcp show adapter`)
- ❖ SANsurfer CLI

Emulex: Download and install the following from the Emulex Web site:

- ❖ Applications kit: Ensure that the application kit version matches the bundled Emulex driver version.

Setting the required HBA and bundled driver parameters

You need to set both HBA and bundled driver parameters to ensure that the multipathing and storage system failover work correctly.

QLogic HBA and bundled driver parameters: To set the bundled HBA parameters, complete the following steps.

Step	Action
1	Run the <code>scli</code> command: <code>scli</code>
2	Set the following HBA parameters: ◆ Port Down Retry Count to 30 ◆ Link Down Timeout to 20
3	Configure the QLogic HBA port topology to Point to Point Only while connecting to a storage system operating in SSI, Standby or Partner CF modes, or to Loop Only if the storage system is operating in Dual Fabric or Mixed CF modes.

Step	Action		
4	Configure the QLogic port speed ideally to the highest speed possible, depending on its ceiling and that of the switch or target port it is connected to. To do this select the Data Rate option and set it to the specified speed.		
5	Select Commit Changes for this HBA so that the changes are applied to the HBA ports.		
6	Repeat the settings for each HBA port listed.		
7	<table> <tr> <td data-bbox="470 526 702 586">If...</td><td data-bbox="702 526 1249 586">Then...</td></tr> </table>	If...	Then...
If...	Then...		

Step	Action
	You are using Red Hat Enterprise Linux 4 - For sanlun fc show adapter command to work, load a QLogic iocli module qioctlmod separately. To load the module, enter the following command: <pre>modprobe -v qioctlmod</pre> - In the modprobe.conf file, append the following line: <pre>options qla2xxx qlport_down_retry=30 ql2xsuspendcount=6</pre> - Re-create the initrd by running the following command: <pre>mkinitrd -f /boot/initrd-"`uname -r`.img `uname -r` -- with=qioctlmod</pre> Note _____ Take a backup of the current initrd before running the mkinitrd command. Taking a backup is a precautionary recommendation. - After rebooting the host, verify that the values of qlport_down_retry and ql2xsuspendcount parameters in the corresponding qla2xxx directory in either /sys/module/qla2xxx/ or /sys/module/qla2xxx/parameters/ match the values in the modprobe.conf file.

Step	Action	
		Example: You can check the value of ql2xsuspendcount on a Red Hat Enterprise Linux host by running the following command: <pre>cat /sys/module/qla2xxx/ql2xsuspendcount</pre>
	You are using Red Hat Enterprise Linux 5	No changes are required for the QLogic bundled driver in the modprobe.conf file, the default settings work fine. Note_____ There are no qla2300/qla2400/qla2500 firmware modules in the Red Hat Enterprise Linux 5 series. These modules have been integrated into the base qla2xxx module itself. _____

Step	Action	
	You are using SUSE Linux Enterprise Server 10	a. In the <code>modprobe.conf.local</code> file, append the following line: <pre>options qla2xxx qlport_down_retry=30</pre> b. Re-create the <code>initrd</code> by running the <code>mkinitrd</code> command. Note_____ Take a backup of the current <code>initrd</code> before running the <code>mkinitrd</code> command. Taking a backup is a precautionary recommendation. c. After rebooting the host, verify that the value of <code>qlport_down_retry</code> parameter in the corresponding <code>qla2xxx</code> directory in <code>/sys/module/qla2xxx/parameters/</code> match the values in the <code>modprobe.conf.local</code>. Example: You can check the value of <code>qlport_down_retry</code> on a SUSE Enterprise Linux Server host by running the following command: <pre>cat /sys/module/qla2xxx/parameters/qlport_down_retry</pre>

Emulex HBA and bundled driver parameters: No changes are required.

Setting the required HBA and bundled driver parameters for Veritas

You need to set both HBA and bundled driver parameters for Veritas to ensure that the multipathing and storage system failover work correctly.

QLogic HBA and bundled driver parameters: To set the bundled HBA parameters, complete the following steps.

Step	Action		
1	Run the scli command: scli		
2	Set the following HBA parameters: ◆ Port Down Retry Count to 30 ◆ Link Down Timeout to 20		
3	Configure the QLogic HBA port topology to Point to Point Only while connecting to a storage system operating in SSI, Standby or Partner CF modes, or to Loop Only if the storage system is operating in Dual Fabric or Mixed CF modes.		
4	Configure the QLogic port speed ideally to the highest speed possible, depending on its ceiling and that of the switch or target port it is connected to. To do this select the Data Rate option and set it to the specified speed.		
5	Select Commit Changes for this HBA so that the changes are applied to the HBA ports.		
6	Repeat the settings for each HBA port listed.		
7	<table border="1"> <tr> <td>If...</td><td>Then...</td></tr> </table>	If...	Then...
If...	Then...		

Step	Action	
	You are using SUSE Linux Enterprise Server 9	a. In the <code>modprobe.conf.local</code> file, append the following line: <pre>options qla2xxx qlport_down_retry=30 ql2xsuspendcount=6 MaxRetriesPerPath=10 MaxRetriesPerIo=10 ql2xfailover=0</pre> b. Re-create the <code>initrd</code> by running the <code>mkinitrd</code> command. Note_____ Take a backup of the current <code>initrd</code> before running the <code>mkinitrd</code> command. Taking a backup is a precautionary recommendation.

Emulex HBA and bundled driver parameters for Veritas: No changes are required.

Note_____

Disable `subfs` and `hwscan` for SUSE Linux Enterprise Server 9 by setting `HOTPLUG_USE_HWSCAN=no` and `HOTPLUG_USE_SUBFS=no` in `/etc/sysconfig/hotplug`.

Recording the WWPNs of the HBAs

About WWPNs

The storage system uses the WWPNs of the HBA ports to identify them. The WWPn is added to an igroup to enable the initiator to access a particular LUN on the storage system.

The WWPn consists of 16 hexadecimal characters. When you enter the characters to create an igroup, separate each pair of characters with a colon. For example, 21:01:00:E0:8B:3A:BB:30.

Each physical port has its own WWPn. A two-port HBA has two unique WWPns.

The FC also defines a worldwide node name (WWNN). The WWNN is not used when creating igroups on the storage system.

Obtaining the WWPn for QLogic HBAs

To obtain the WWPns of QLogic HBAs by using the SANsurfer FC HBA CLI client, complete the following steps.

Step	Action
1	Enter the following command on the Linux host console: <code>scli -i</code>
2	Record the WWPn for each port listed.

For more information about using the SANsurfer FC HBA CLI client, see the *SANsurfer FCP HBA CLI Application User's Guide*. To download the guide, go to http://support.qlogic.com/support/drivers_software.asp and then click your HBA model.

The WWPns are also printed on a sticker at the bottom of most QLogic HBAs.

Obtaining the WWPN for Emulex HBAs

To obtain the WWPNs for using the Emulex HBAnyware for Linux CLI, complete the following steps.

Step	Action
1	Enter the following command on the Linux host console: <code>/usr/sbin/hbanyware/hbacmd ListHBAs</code>
2	Record the WWPN for each port listed.

For more information about Emulex HBAnyware for Linux CLI, see the *Emulex Driver and Utilities for Linux User Manual*. To download the guide, go to <http://www.emulex.com/ts/index.html>, click NetApp, click your HBA model, and then click Linux.

Configuring the Storage System for FC

Before you begin

You need the host's WWPNs that you recorded in Chapter 2, "[Installing HBAs and Software on the Host](#)," on page 15, before you can configure the storage system. You use the WWPNs when creating igroups on the storage system.

Tasks for configuring the storage system

To configure the storage system and prepare it for initiator access, you must complete the following tasks:

- ◆ Ensure that the FC service is licensed and started.
- ◆ For active/active storage system configurations, verify that the storage system is running a supported cfmode by running the `fcpx show cfmode` command. The support matrix lists the supported cfmode values for your configuration. If you need to change the cfmode, see Chapter 4, "[Migrating a Storage System to single_image mode](#)," on page 39.

For standby and mixed cfmode, special Fibre Channel switch zoning is required. See "[Zoning with standby or mixed cfmode](#)" on page 48.

- ◆ Create LUNs, create an igroup using the host's WWPNs, and then map the LUNs to the igroup for the Linux host. You can manage LUNs using the Web-based FilerView interface or using the Data ONTAP command line. Be sure to specify the LUN type and FC igroup type as `linux`, and be sure a LUN is mapped as LUN 0.
- ◆ **Optional:** Enable ALUA, if you have not already done so.

Note

ALUA is not supported with Veritas Storage Foundation.

- ◆ If you are using clustered Linux hosts, the igroup must contain the WWPNs of all hosts in the cluster that need access to the mapped LUN.

For detailed configuration steps

For detailed storage system configuration steps, see the *Block Access Management Guide* for your version of Data ONTAP.

About this chapter

This chapter describes how to migrate an active/active storage system configuration to single_image mode.

You can migrate to single_image mode while installing the new FC Linux Host Utilities, or as a stand-alone project. If you are migrating during Host Utilities installation, you should make the storage system changes after the old Linux HBA drivers have been removed, and before connecting the Fibre Channel cables to the Linux host HBAs.

Using single_image mode requires dm-multipath or Veritas DMP on the Linux host.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Understanding cfmode](#)” on page 40
- ◆ “[Migrating to single_image mode](#)” on page 42

Understanding cfmode

What cfmode is

Cluster failover mode (cfmode) is a storage system FC failover mode setting which controls how the appliance handles Fibre Channel before and after cluster failover events. The cfmode is not available and not relevant on non-clustered storage controllers. NetApp supports the following cfmodes.

- ◆ Single System Image (single_image) cfmode
- ◆ Dual Fabric (dual_fabric) cfmode
- ◆ Partner cfmode
- ◆ Standby cfmode
- ◆ Mixed cfmode

In this section, only the single_image cfmode is explained in detail.

Impact of changing the cfmode setting

Changing the cfmode setting of an active/active storage affects the following items:

- ◆ Access to LUNs
When you change the cfmode setting, you change the available paths between the host and the storage system. Some previously available paths are no longer available and some new paths become available. You must reconfigure every host that is connected to the storage system to discover the new paths. The LUNs might be accessible but must not be used until you reconfigure the host.
- ◆ Multipathing policy
single_image mode requires dm-multipath or Veritas DMP on a Linux host. Changing the cfmode setting might also affect the multipathing policy because there are more paths between the host and the storage system.
- ◆ Zoning
Updating the storage system to single_image cfmode changes the WWNN of the storage system and the WWPNs of the target ports on each controller. If the storage system connects to a switch that is zoned by WWPN, you must check the zoning setup and make the required updates to your switch. Switches that are zoned by physical port are not affected.
- ◆ Switch configuration and cabling
In mixed and dual fabric modes, the target ports log into the fabric in loop mode. In single_image mode, the target ports log into the fabric using point-

to-point mode. This means you might need to update your switch configuration when you migrate to `single_image` mode so that the target connects to the fabric as a point-to-point device.

Any currently supported cabling configuration is supported by `single_image` mode. This means that if you have a supported configuration, it will continue to work after you migrate to `single_image` mode. To optimize high availability and performance, you can change the cabling from the storage system to the fabric.

Note

For detailed information about the cabling configurations supported, and recommended examples for systems with onboard FC ports, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Prerequisites for changing cfmode

Complete the following tasks before you change the `cfmode`.

◆ **Save configuration information:**

Enter commands that display information about the current configuration and save the output to the local disk.

On the host, make a backup of all the volume groups that contain LUNs. These backups must contain up-to-date information for both the data within the volume group and the volume group configuration. Perform a complete backup of the system's data.

Note

If you are using a native disk manager, then use the appropriate commands to obtain information about your disks.

◆ **Identify affected hosts and initiators:**

You must reconfigure and reboot any hosts and initiators that access LUNs on the systems you are upgrading.

Enter the following command:

igroup show

◆ **Schedule downtime for the storage system:**

Hosts cannot access LUNs when you are changing `cfmode`. You must allow sufficient downtime to change the `cfmode` setting, and reconfigure and reboot the host. Enter commands that quiesce host I/O and take host applications off-line.

Migrating to single_image mode

What single_image mode is

In single_image mode, both controllers in an active/active storage system configuration share the same WWNN. Single_image mode provides the following advantages:

- ◆ The host can access all LUNs through any target port on the active/active configuration.
- ◆ The single_image mode is supported on all storage systems and by all supported FC hosts beginning with Data ONTAP 7.1.
- ◆ There are no switch limitations. You can connect a storage system in single_image mode to any supported FC switch.

Reasons for using single_image mode

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

- ◆ To increase the number of paths to the active/active configuration
You can increase the number of available paths in your configuration because all ports on each controller are available to the host.
- ◆ To support heterogeneous configurations

To learn more about how single_image mode works, see the *Block Access Management Guide for iSCSI and FCP* for your version of Data ONTAP; however, follow the procedures and recommendations in this document to migrate to single_image mode.

Resolving LUN mapping conflicts

In single_image mode, make sure the following mapping conflicts do not occur:

- ◆ Two LUNs, each one on a different controller in the active/active configuration, that have the same LUN ID and are mapped to the same igroup.
- ◆ Initiators that are members of igroups with different operating system types (for example, an initiator that is a member of a Linux igroup on one controller and a member of an AIX® igroup on another controller).

When you change the cfmode and restart the FC service, one WWNN is automatically assigned to both controllers in the active/active storage system configuration. This means you can have LUN mapping conflicts between the systems in the active/active configuration.

Before you switch to the `single_image` `cfmode` setting, you use the `lun config_check -S` command to check for any LUN mapping and operating system type conflicts in the active/active configuration.

Tasks for migrating to single_image mode

To migrate an FC-attached host and storage system to `single_image` mode you must complete the following tasks:

1. Prepare the host: Quiesce I/O and protect the configuration by unmounting file systems and disconnecting LUNs attached to the host. Configure `dm-multipath` if not already running.
2. Upgrade the storage system to `single_image` mode: Resolve LUN mapping or OS-type conflicts, and change the `cfmode` setting.
3. Validate zoning and cabling: Check zoning and resolve zone mapping issues on configurations that use soft zoning.
4. Reconfigure the host: Reconfigure the host to detect the new devices and new paths to the devices.

Preparing the host

To prepare the Linux host for migrating to `single_image` mode, complete the following steps.

Step	Action
1	Quiesce host I/O and stop any applications accessing LUNs on the storage system.
2	Disconnect the host by unmounting file systems on LUNs, and then unloading the drivers for the Fibre Channel HBAs.
3	Proceed to “ Migrating to single_image mode ” on page 44.

Migrating to single_image mode

To migrate the storage system to single_image mode, complete the following steps.

Step	Action
1	Open a window to the console of each storage system controller.
2	Log in at the console of the systems you are upgrading, and enter the following command to identify and resolve LUN and igroup mapping conflicts: <pre>lun config_check -S</pre> Result: Mappings that are displayed using the asterisk character (*) indicate an error, and also provide LUN and igroup information that you can use to resolve conflicts. Use the lun map command to re-map LUNs to a new ID that is unique across the storage system. You use the igroup command to change igroup operating system types, if required.
3	Switch to advanced privileges by entering the following command: <pre>priv set advanced</pre>
4	Stop the FC service on both systems. On each system, enter the following command: <pre>fcp stop</pre>
5	Set the cfmode to single_image. On each system console, enter the following command: <pre>fcp set cfmode single_image</pre> If you see messages indicating mapping conflicts, then return to Step 2 to resolve the mapping conflicts.
6	Start the FC service on each system in the active/active configuration. On each system, enter the following command: <pre>fcp start</pre>
7	After you switch to the single_image setting, both controllers in the active/active configuration use the same WWNN. Obtain the WWNN of the storage system by entering the following command on the console of each system: <pre>fcp nodename</pre>

Step	Action
8	If you are using switch zoning by WWPN, you need the WWPNs of the target ports to reconfigure zones on your switch. To obtain the WWPNs, enter the following command on each controller in the active/active configuration: <pre>fcp config</pre> For more information about zoning, see “Validating zoning and connectivity” on page 45.
9	Return to administrative privileges. On each storage system, enter the following command: <pre>priv set admin</pre>

Validating zoning and connectivity

Changes to WWNNs and WWPNs: After you change the cfmode setting to `single_image`, both controllers share a single WWNN. (One controller in the active/active configuration takes the WWNN of the other.)

The WWPNs of the target ports also change. You might have to update your zoning configuration depending on the type of zoning you use. If you zone by WWPN, you must update your zones to use the new WWPNs. Configurations that are zoned by physical port are not affected.

Cabling and zoning: If you do not change your cabling, migrating to `single_image` mode changes the number of available paths to the LUNs. Additional paths to the target become available because more target ports are visible to the host. You can use zoning or Data ONTAP port sets to reduce the number of paths in your configuration to comply with the number of paths supported by your multipathing software.

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

For details about port sets, see the *Block Access Management Guide for iSCSI and FCP* for your version of Data ONTAP.

You can also change your cabling to reduce the number of paths. It is recommended that after you change the cfmode setting, make sure each HBA has one primary and at least one secondary path to each LUN. For additional information, see the NetApp Interoperability Matrix at

<http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Reconfiguring the hosts

To reconfigure Linux hosts, reload the FC external HBA drivers and mount file systems. Verify that all LUNs were re-discovered by running the sanlun utility.

Configuring FC Cabling and Switches

Fabric or direct connections

You can connect your host to the storage system using either Fibre Channel switches or direct cables. A switched fabric provides more flexibility and enables you to connect more hosts. A direct connected configuration requires one target port on each storage system controller for each host.

Before connecting the Linux host to Fibre Channel switches

Do not connect the Linux host to the Fibre Channel switches until at least one LUN has been created and mapped to the host. Be sure to connect the storage system and zone the switches before connecting the host.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Switched fabric configurations”](#) on page 48
- ◆ [“Direct-connected configurations”](#) on page 50

Switched fabric configurations

High availability

To avoid single points of failure, use two physically independent Fibre Channel switch fabrics to connect your Linux host to an active/active storage system configuration.

Configuration and zoning details

For details about supported Fibre Channel configurations and switch zoning, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Use default HBA and target topology

For a fabric-connected configuration, the default autonegotiated HBA and storage system target Fibre Channel topology works correctly. Be sure you have set the HBA and driver parameters to the values specified in “[Setting the required HBA and driver parameters](#)” on page 24.

Single initiator zoning

When using the switched fabric, it is always recommended to use single initiator zoning for each of the initiator ports, that is, each zone should have only a single initiator port.

Zoning with standby or mixed cfmode

If the active/active configuration storage system is using the standby or mixed cfmode, your switches must be zoned by WWPn.

To prevent failover problems, create zones that include the host HBAs and the primary port WWPns on the storage system. Do not include the standby port WWPns on the storage system in the zones with the host HBA ports.

The host should never access the standby port WWPns. During failover, the standby ports of the takeover system assume the WWPns of the primary ports of the system that is down. The primary ports are already included in existing zones, so the host can access them.

Connect cables

Connect the storage system cables to the Fibre Channel switches first. After you have created LUN 0 and mapped it to the host, you can connect the Fibre Channel cables from the host HBAs to the Fibre Channel switches.

Direct-connected configurations

High availability

To avoid single points of failure, use two Fibre Channel cables to connect your Linux host to both controllers of an active/active storage system configuration.

Use default HBA and target topology

For a direct connected configuration, the default autonegotiated HBA and storage system target Fibre Channel topology works correctly. Be sure you have set the HBA and driver parameters to the values specified in “[Setting the required HBA and driver parameters](#)” on page 24.

Single_image cfmode required

A direct Fibre Channel connection to an active/active storage system configuration requires that the storage system be using single_image cfmode.

To verify cfmode, enter the following command on the storage system console:

```
fcpx show cfmode
```

For information about changing to single_image cfmode, see Chapter 4, “[Migrating a Storage System to single_image mode](#),” on page 39.

Connect cables

After you have created LUN 0 and mapped it to the host, you can connect the Fibre Channel cables from the host to the target ports on the storage system.

About this chapter

This chapter applies only to Red Hat Enterprise Linux hosts. For SUSE Linux Enterprise Server hosts, see Chapter 7, “[Accessing LUNs from a SUSE Linux Enterprise Server Host](#),” on page 83.

If you are using the dm-multipath support or Veritas DMP, the method you use to access LUNs is different from when you are not using multipathing. This chapter has three main sections, one for accessing LUNs with dm-multipath, one for accessing LUNs with Veritas DMP, and one for accessing LUNs without dm-multipath.

Note

For information about the supported configuration, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Loading the external HBA driver on the host”](#) on page 53
- ◆ [“Loading the bundled HBA driver on the host”](#) on page 54
- ◆ [“Configuring dm-multipath”](#) on page 55
- ◆ [“Starting multipath service”](#) on page 61
- ◆ [“Accessing LUNs with dm-multipath support”](#) on page 62
- ◆ [“Stopping dm-multipath”](#) on page 71
- ◆ [“Accessing LUNs with Veritas DMP support”](#) on page 72
- ◆ [“Accessing LUNs without dm-multipath support”](#) on page 73
- ◆ [“Unloading the external HBA driver and shutting down”](#) on page 81
- ◆ [“Unloading the bundled HBA driver and shutting down”](#) on page 82

Loading the external HBA driver on the host

Loading the external HBA driver

Once you have at least one LUN mapped to the Linux host, you can load the external HBA driver. When the external HBA driver is started, it scans the storage systems and discovers all mapped LUNs. If you configure dm-multipath or Veritas DMP support, multipath devices are created for all LUNs discovered.

Before you load the driver

Be sure that you have at least one LUN mapped to the Linux host as LUN 0 before loading the external HBA driver.

Starting FC

To start the FC on a Linux host, load the driver for your HBAs.
Be sure you have configured multipathing before starting FC.

Loading the external QLogic HBA driver: Complete the following step to load the external QLogic HBA driver.

Step	Action
1	Enter the following command on the Linux host: <code>modprobe -v qla2300</code> <code>modprobe -v qla2400</code> <code>modprobe -v qla2500</code>

Loading the external Emulex HBA driver: Complete the following step to load the external Emulex HBA driver.

Step	Action
1	Enter the following command on the Linux host: <code>modprobe -v lpfc</code>

Loading the bundled HBA driver on the host

Loading the bundled HBA driver

Once you have at least one LUN mapped to the Linux host, you can load the bundled HBA driver. When the bundled HBA driver is started, it scans the storage systems and discovers all mapped LUNs. If you configure dm-multipath or Veritas DMP support, multipath devices are created for all LUNs discovered.

Before you load the bundled driver

Be sure that you have at least one LUN mapped to the Linux host as LUN 0 before loading the bundled HBA driver.

Starting FC

To start the FC on a Linux host, load the bundled driver for your HBAs.

Be sure you have configured multipathing before starting FC.

Loading the bundled QLogic HBA driver: Complete the following step to load the bundled QLogic HBA driver.

Step	Action	
1	Enter the following command on the Linux host:	
	If...	Then...
	You are using Red Hat Enterprise Linux 4	<code>modprobe -v qla2300</code> <code>modprobe -v qla2400</code> <code>modprobe -v qla2500</code>
	You are using Red Hat Enterprise Linux 5	<code>modprobe -v qla2xxx</code>

Loading the bundled Emulex HBA driver: Complete the following step to load the bundled Emulex HBA driver.

Step	Action
1	Enter the following command on the Linux host: <code>modprobe -v lpfc</code>

Configuring dm-multipath

Note—
Skip this section if you are configuring LUNs without dm-multipath support (“[Accessing LUNs without dm-multipath support](#)” on page 73) or if you are using Veritas Storage Foundation (“[Accessing LUNs with Veritas DMP support](#)” on page 72).

Editing the configuration file

To configure dm-multipath, complete the following steps.

Step	Action	
1	Check if the /etc/multipath.conf file exists.	
	If...	Then...
	The file exists	a. Open the file. b. Comment out all the lines. c. Append the sample specified in “Red Hat Enterprise Linux 5.2” on page 57 or see the note on “page 58”, depending on your version of the operating system.
	The file does not exist	Create a new file with the entries as specified“ Red Hat Enterprise Linux 5.2 ” on page 57 or see the note on “ page 59 ”, depending on your version of the operating system.
2	Save the changes.	

Step	Action
3	Edit the <code>/etc/multipath.conf</code> file to exclude (blacklist) all devices that do not correspond to any LUN configured on storage controller and which are mapped to your Linux host (that is, blacklist those devices that are not displayed in the <code>sanlun lun show</code> command). Complete the following steps to blacklist the devices. - Blacklist the non NetApp SCSI devices installed on your host, using the WWID method. To obtain the WWID of a device, run the <code>scsi_id</code> command on the respective device. Example: Suppose <code>/dev/sda</code> is a local SCSI drive. To obtain the WWID, run the following command: <pre>scsi_id -gus /block/sda</pre> The output looks like as specified below: <pre>SIBM-ESXSMAW3073NC_FDAR9P66067WJ</pre> Add this entry for blacklisting. - Blacklist other devices using the devnode method by adding the following: <pre>devnode "^hd[a-z]"</pre> <pre>devnode "^(ram raw loop fd md dm- sr scd st)[0-9]*"</pre> <pre>devnode "^cciss!c[0-9]d[0-9]*"</pre> Example: On a Red Hat Enterprise Linux 4 host: <pre>devnode_blacklist { devnode "^(ram raw loop fd md dm- sr scd st)[0-9]*" devnode "^hd[a-z][0-9]*" devnode "^cciss!c[0-9]d[0-9]*p[0-9]*" wwid SIBM-ESXSMAW3073NC_FDAR9P66067WJ }</pre>

Step	Action
4	If you are using ALUA, specify the ALUA callout program, that is, the <code>/sbin/mpath_prio_alua</code> program, instead of <code>/sbin/mpath_prio_ontap</code> . ALUA callout programs are available in the Red Hat Enterprise Linux distribution.
5	Save the changes.

The following example shows the required changes to the `/etc/multipath.conf` file.

Red Hat Enterprise Linux 5.2:

```
defaults
{
    user_friendly_names    yes
    max_fds                4096
}
blacklist
{
    wwid <DevId>
    devnode "^ (ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z]"
    devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
}
devices
{
    device
    {
        vendor                "NETAPP"
        product               "LUN"
        getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
        prio_callout           "/sbin/mpath_prio_ontap /dev/%n"
        features               "1 queue_if_no_path"
        hardware_handler       "0"
        path_grouping_policy   group_by_prio
        failback               immediate
        rr_weight              uniform
        rr_min_io              128
        path_checker            directio
    }
}
```

<DevID> is the unique ID. See Step 3 of “[Editing the configuration file](#)” on page 55 to generate wwid.

Note

directio pathchecker is available with Red Hat Enterprise Linux 5 Update 1 and later. For Red Hat Enterprise Linux 5.0, use readsector0 pathchecker.

max_fds is available with Red Hat Enterprise Linux 5 Update 2 and later. For Red Hat Enterprise Linux 5.0 and Red Hat Enterprise Linux 5 Update 1, max_fds is not supported.

Red Hat Enterprise Linux 4.7:

```
defaults
{
    user_friendly_names    yes
    rr_min_io              128
    max_fds                4096
}
devnode_blacklist
{
    wwid <DevId>
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z]"
    devnode "^cciss!c[0-9]d[0-9]*[p[0-9]]*"
}
devices
{
    device
    {
        vendor              "NETAPP"
        product             "LUN"
        flush_on_last_del   yes
        getuid_callout       "/sbin/scsi_id -g -u -s /block/%n"
        prio_callout         "/sbin/mpath_prio_ontap /dev/%n"
        features            "1 queue_if_no_path"
        hardware_handler     "0"
        path_grouping_policy group_by_prio
        failback             immediate
        rr_weight            uniform
        path_checker         directio
    }
}
```

<DevID> is the unique ID. See Step 3 of “[Editing the configuration file](#)” on page 55 to generate wwid.

Note

directio pathchecker is available with Red Hat Enterprise Linux 4 Update 7 and later. For Red Hat Enterprise Linux 4 Update 6 and earlier, use readsector0 pathchecker.

flush_on_last_del is available with Red Hat Enterprise Linux 4 Update 7 and later.

max_fds is available with Red Hat Enterprise Linux 5 Update 2 and later. For Red Hat Enterprise Linux 5.0 and Red Hat Enterprise Linux 5 Update 1, max_fds is not supported.

For more information about tuning the multipath configuration, see “[Tuning the dm-multipath configuration](#)” on page 70.

Default versus recommended settings

The default versus recommended settings for Red Hat Enterprise Linux 4 and Red Hat Enterprise Linux 5 are as follows:

Component	Parameter	Red Hat Enterprise Linux 4		Red Hat Enterprise Linux 5	
		Default value	NetApp recommended value	Default value	NetApp recommended value
Qlogic External Driver	MaxRetriesPerIo	0	10	0	10
	MaxRetriesPerPath	0	10	0	10
	ql2xsuspendcount	10	6	10	6
Qlogic Bundled Driver	qlport_down_retry	0	30	30	30
	ql2xsuspendcount	10	6	NA	NA
Qlogic FC HBA	Link Down Timeout	30	20	30	20
Multipath.conf	user_friendly_names	no	yes	no	yes
Veritas DMP	dmp_restore_interval	300	60	300	60

Note

On Red Hat Enterprise Linux 4, the `user_friendly_names` feature is supported only from Red Hat Enterprise Linux 4 Update 5 onwards.

On Red Hat Enterprise Linux 5, the `user_friendly_names` feature is disabled by default. However, it is enabled by default in the Red Hat Enterprise Linux 5 Anaconda installer for enabling the SAN boot root on dm-multipath feature. For Emulex, the default settings work fine, that is, there are no recommended settings for Emulex.

Starting multipath service

Note—Skip this section if you are configuring LUNs without dm-multipath support (“[Accessing LUNs without dm-multipath support](#)” on page 73) or if you are using Veritas Storage Foundation (“[Accessing LUNs with Veritas DMP support](#)” on page 72).

Once you have loaded the external HBA drivers, start the multipath service on the Linux host. Complete the following steps.

Step	Action
1.	To start the multipath daemon, enter the following command: <code>/etc/init.d/multipathd start</code>
2.	To configure the dm-multipath devices, run the following command: <code>multipath</code>

Configuring the multipath service to start automatically after reboot: Once you have started the multipath service on the Linux host, you can configure it to start automatically after reboot. Complete the following step.

Step	Action
1.	To add the multipath service to the boot sequence, enter the following commands on the Linux console: <code>chkconfig --add multipathd</code> <code>chkconfig multipathd on</code>

Accessing LUNs with dm-multipath support

Note

Skip this section if you are configuring LUNs without dm-multipath support ([“Accessing LUNs without dm-multipath support”](#) on page 73) or if you are using Veritas Storage Foundation ([“Accessing LUNs with Veritas DMP support”](#) on page 72).

To access LUNs with dm-multipath support, perform the following tasks.

1. Verify the HBA setting and dm-multipath configuration. For more information, see [“Verifying the HBA setting and dm-multipath configuration”](#) on page 63.
2. Identify and view the list of LUNs. For more information, see [“Identifying and viewing a list of LUNs”](#) on page 66.
3. Create and mount the file system. For more information, see [“Creating and mounting a file system on a multipath device”](#) on page 68.
4. Tune the multipath configuration if required. For more information, see [“Tuning the dm-multipath configuration”](#) on page 70.

Verifying the HBA setting and dm-multipath configuration

Overview

Once you start dm-multipath, you should verify that the HBA settings are correct, the desired multipathing is correctly configured, and the multipath maps were created.

Verifying the HBA settings

QLogic:

Verify the ql2xsuspendcount, MaxRetriesPerIo, and MaxRetriesPerPath parameters in the following way:

Check the values for these parameters in the corresponding qla2xxx directory in the sysfs, /sys/module/qla2xxx/ or /sys/module/qla2xxx/parameters/.

Example: You can check the value of ql2xsuspendcount on a Red Hat Enterprise Linux host by running the following command:

```
cat /sys/module/qla2xxx/ql2xsuspendcount
```

PortDownRetryCount and LinkDownTimeout along with Connection Options and Data Rate (which need to be manually set instead of the default autonegotiate) can be verified as follows.

Run the **-c all** option of the **scli** command:

```
/opt/QLogic_Corporation/SANsurferCLI/scli -c all
```

Example:

The following is a sample output of the **scli** command:

```
-----
HBA 0: QLE2462 Port 1 WWPN 21-00-00-E0-8B-80-72-A1 Port ID 03-05-00
-----
Connection Options : 1 - Point-to-Point Only.
Data Rate : 4 Gbps
Frame Size : 2048
Hard Loop ID : 0
Loop Reset Delay (seconds) : 5
Enable Host HBA BIOS : Disabled
Enable Hard Loop ID : Disabled
Enable FC Tape Support : Enabled
Operation Mode : 0 - Interrupt for every I/O completion
```

```

Interrupt Delay Timer (100ms) : 0
Execution Throttle : 16
Login Retry Count : 8
Port Down Retry Count : 30
Enable LIP Full Login : Enabled
Link Down Timeout (seconds) : 20
Enable Target Reset : Enabled
LUNs Per Target : 128
-----
HBA 1: QLE2462 Port 2 WWPN 21-01-00-E0-8B-A0-72-A1 Port ID 03-01-00
-----
Connection Options : 1 - Point-to-Point Only.
Data Rate : 4 Gbps
Frame Size : 2048
Hard Loop ID : 0
Loop Reset Delay (seconds) : 5
Enable Host HBA BIOS : Disabled
Enable Hard Loop ID : Disabled
Enable FC Tape Support : Enabled
Operation Mode : 0 - Interrupt for every I/O completion
Interrupt Delay Timer (100ms) : 0
Execution Throttle : 16
Login Retry Count : 8
Port Down Retry Count : 30
Enable LIP Full Login : Enabled
Link Down Timeout (seconds): 20
Enable Target Reset : Enabled
LUNs Per Target : 128

```

Emulex:

The default settings work fine in case of Emulex.

Verifying the configuration

Use the **multipath** command on the Linux host to view the dm-multipath configuration. You can change the amount of detail displayed by using the **-v** option. For details on the multipath command options, enter **multipath -h**.

To view all the configuration details, enter the following command on the Linux host:

```
multipath -v3 -d -ll
```

Note that the **-d** (“dry run”) parameter prevents the command from actually updating the multipath maps.

Verify that the multipathd daemon is running by issuing the following command:

```
/etc/init.d/multipathd status
```

Viewing the multipath devices: To view a list of multipath devices, including which /dev/sdx devices are used, enter the following command:

```
multipath -ll
```

Example: The following output shows two LUNs, with four paths each—two proxy and two non-proxy paths.

```
360a98000486e5372635a44646a505643NETAPP, LUN
[size=10G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=8][active]
  \_ 6:0:0:0 sda 8:0 [active][ready]
  \_ 7:0:1:0 sdg 8:96 [active][ready]
\_ round-robin 0 [prio=2][enabled]
  \_ 6:0:1:0 sdc 8:32 [active][ready]
  \_ 7:0:0:0 sde 8:64 [active][ready]
360a98000486e5363693444646a2f656cNETAPP, LUN
[size=10G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=8][enabled]
  \_ 6:0:1:1 sdd 8:48 [active][ready]
  \_ 7:0:0:1 sdf 8:80 [active][ready]
\_ round-robin 0 [prio=2][enabled]
  \_ 6:0:0:1 sdb 8:16 [active][ready]
  \_ 7:0:1:1 sdh 8:112 [active][ready]
```

Note that the /dev/mapper devices are persistent across reboots, but the /dev/sdx devices are not. After a reboot, or a reload of the external HBA driver, you might find that different /dev/sdx devices make up a specific /dev/mapper device. The /dev/mapper device, however, always corresponds to the same LUN.

To view the /dev/mapper devices, enter the following command:

```
ls -l /dev/mapper/
```

Example: The following example shows sample output of **ls -l dev/mapper** command.

```
total 0
brw----- 1 root root 253, 1 Sep 20 17:09
360a98000486e5363693444646a2f656c
brw----- 1 root root 253, 0 Sep 20 17:09
360a98000486e5372635a44646a505643
lrwxrwxrwx 1 root root    16 Sep 12 10:16 control -> ../device-
mapper
```

Identifying and viewing a list of LUNs

Persistently identifying LUNs

The dm-multipath code creates persistent devices for each LUN in the `/dev/mapper/` directory on the Linux host. Use these entries for multipath LUNs. You can create a file system directly on a multipath device in `/dev/mapper/`.

The `/dev/mapper/` entries can also be used in the `/etc/fstab` file.

Viewing list of LUNs

To view information about LUNs, use the `sanlun` command.

Using `sanlun`: To view a list of LUNs that are mapped to the Linux host, enter the following command on the Linux host console:

```
sanlun lun show all
```

Example: The following output shows one LUN with four paths mapped to the Linux host.

```
host:/ # sanlun lun show all
filer: lun-pathname device filename adapter protocol lun size lun state
f3020-92: /vol/vol24/lun24 /dev/sdt host4 FCP 5g (5368709120) GOOD
f3020-93: /vol/vol24/lun24 /dev/sdu host4 FCP 5g (5368709120) GOOD
f3020-93: /vol/vol24/lun24 /dev/sdk host3 FCP 5g (5368709120) GOOD
f3020-92: /vol/vol24/lun24 /dev/sdj host3 FCP 5g (5368709120) GOOD
```

To display multipath information for native device-mapper multipath, enter the following command:

```
sanlun lun show -p
```

Example: The following output shows when device-mapper multipathing is configured.

```

[root@lnx199-110 ~]# sanlun lun show -p
f3020-92:/vol/vol1/lun2 (LUN 1)          Lun state: GOOD
Lun Size:      3.6g (3879731200)    Filer_CF_State: Cluster Enabled
Protocol: FCP          Filer Partner: f3020-93
DM-MP DevName: mpath3    (360a9800043346852563444717a513571)    dm-1
Multipath-provider: NATIVE
-----
sanlun      Filer      Primary      Partner
path        Path    /dev/      Host        Filer      Filer
state       type   node      HBA         port      port
-----
UP  primary      sde      host17      1b         --
UP  primary      sdg      host18      1a         --
UP  secondary    sdc      host17      --         1a
UP  secondary    sdi      host18      --         1b

```

Creating and mounting a file system on a multipath device

Overview

You can optionally create a file system on the multipath device that represents a LUN.

When creating a file system, you use the multipath device in the `/dev/mapper/` directory.

Creating the file system

To create a file system on a LUN, enter the following command on the Linux host console:

```
mkfs -t type /dev/mapper/device
```

type is the file system type, such as `ext3`.

device is the multipath device name of the LUN in the `/dev/mpath/` directory.

Example: The following command creates an `ext3` file system on a dm-multipath device:

```
mkfs -t ext3 /dev/mapper/360a980004334686267344374474d4f75
```

Mounting the file system

Multipath devices in the `/dev/mapper/` directory that are listed in `/etc/fstab` are automatically remounted after a reboot, assuming you have correctly configured the multipath service to start automatically.

To mount the new file system, you create a mountpoint, and add an entry to the `/etc/fstab` file. To mount the file system, complete the following steps.

Step	Action
1	Create a mountpoint using the following command: <pre>mkdir <i>mount_point</i></pre> <i>mount_point</i> is the name of the mountpoint you are creating.
2	Open the <code>/etc/fstab</code> file with a text editor.

Step	Action
3	Add an entry to <code>/etc/fstab</code> as follows, separating each entry with Tab (except the final two zeros, which are separated by a space): <pre>device mount_point type defaults 0 0</pre> <i>device</i> is the name of the device in the <code>/dev/mpath/</code> directory. <i>mount_point</i> is the mountpoint you created for the file system. <i>type</i> is the file system type, such as <code>ext2</code> or <code>ext3</code>. Example: <pre>/dev/mapper/360a980004334686267344374474d4f75 /mnt/ssl/lun_0 ext3 defaults 0 0</pre>
4	To mount the file system, enter the following command: <pre>mount mountpoint</pre> Example: <pre>mount /mnt/ssl/lun_0</pre>
5	Do either of the following: ◆ Enter df to display the mounted file systems. Verify the new file system is listed. ◆ Change to the mountpoint directory and try writing and reading a file to verify the file system mounted successfully.

Tuning the dm-multipath configuration

Tuning multipath parameters

Advanced administrators who are familiar with Linux performance tuning might want to modify two of the Linux parameters for best performance.

queue_depth: The queue_depth value for the SCSI device queue is one of the critical settings to obtain the best performance. This value can be found in sysfs, and you can echo a new value in to change this for each device SCSI command queue. The following example shows how to change the SCSI device queue depth for device sdb.

```
# cat /sys/block/sdb/device/queue_depth
32
# echo "64" > /sys/block/sdb/device/queue_depth
# cat /sys/block/sdb/device/queue_depth
64
```

rr_min_io: The rr_min_io setting for dm-multipath specifies the number of I/O operations sent through a path before switching to the next path. Lowering this value from the default value of 1000 has been seen to dramatically affect overall throughput for dm-multipath, especially for large I/O workloads (64 KB or more). Other performance tests have shown a good value for rr_min_io to be N, where N is less than or equal to the number of threads in the system that issue a single I/O and block, waiting for it to complete. Tuning the rr_min_io value is suggested for optimal performance because it really depends on the workload and other factors such as the host memory and CPU speed.

The rr_min_io can be set in the device section of the multipath.conf file as follows:

```
device
{
    rr_min_io      128
}
```

See [“Configuring dm-multipath”](#) on page 55 for details on blacklisting the local hard drives and other resources that should not be included in the multipathing configuration.

Stopping dm-multipath

Note

Skip this section if you are configuring LUNs without dm-multipath support (“[Accessing LUNs without dm-multipath support](#)” on page 73) or if you are using Veritas Storage Foundation (“[Accessing LUNs with Veritas DMP support](#)” on page 72).

When you want to stop the dm-multipath support on the Linux host, you should stop the affected services in the order shown in this section.

To stop dm-multipath and FC on the host, complete the following steps.

Step	Action
1	Unmount all file systems on the FC LUNs.
2	Flush the multipath maps by entering the following command: multipath -F
3	Stop the multipath daemon by entering the following command: /etc/init.d/multipathd stop

Accessing LUNs with Veritas DMP support

Veritas Storage Foundation

If you are using Veritas Storage Foundation, use the Veritas DMP solution. For details about configuring the Veritas Volume Manager and accessing LUNs, see the *Veritas® Volume Manager Administrator's Guide* for Linux that is shipped along with the software.

The `vxddmp restore` daemon should be configured with an interval of 60.

To configure the daemon, run the following command:

```
vxddmptune dmp_restore_interval 60
```

On reboot this value comes into effect and remains persistent.

You can also use the following commands to configure the daemon value.

```
vxddmpadm stop restore
```

```
vxddmpadm start restore interval=60
```

This value comes into effect immediately and is not persistent on reboot.

To verify that the daemon value has been successfully configured, run the following command:

```
vxddmpadm stat restored
```

Accessing LUNs without dm-multipath support

Overview

Without multipathing software, the host has no way of distinguishing multiple LUNs from multiple paths to the same LUN.

If you do not use dm-multipath or Veritas DMP, you must limit each LUN to a single path. If you are using an active/active configuration storage system, you must use the standby cfmode to avoid having a secondary path visible on the storage system.

If you are using a switched Fibre Channel fabric, you must zone your switch to prevent the host from having multiple paths to any one LUN. See “[Zoning with standby or mixed cfmode](#)” on page 48. You can create a switch zone with a single host HBA WWPN and a single storage system WWPN; this zone provides a single path to a LUN.

There are different HBA parameter settings for configurations without dm-multipath. See “[Setting the required HBA and driver parameters](#)” on page 24.

Configurations without dm-multipath can have additional restrictions on host OS versions and HBA models. See the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Tasks for accessing LUNs

After the initiator itself is configured, you need to set up the host to access the LUNs you created on the storage system. To access LUNs, you must complete the following tasks.

1. Identify and view the list of LUNs. For more information, see “[Identifying and viewing LUN information](#)” on page 75.
2. Create one or more partitions of the LUN. A file system needs a partition, even if the one partition uses all the space in the LUN. For more information, see “[Creating a partition](#)” on page 76.
3. Create a file system on the partition. For more information, see “[Creating a file system](#)” on page 78.

4. Create a label for the file system and add an entry for the file system's label in `/etc/fstab`. Mount the file system. For more information, see [“Labeling and mounting the file system”](#) on page 79.

Accessing LUNs without dm-multipath

Identifying and viewing LUN information

**Persistently
identifying the LUN**

LUNs on the storage system are seen by the Linux host as SCSI devices. Linux assigns a unique device name, such as /dev/sda or /dev/sdb, to each LUN. The device names are assigned as the LUNs are discovered. Because of variable network delays and other issues, the order in which LUNs are discovered cannot be predicted. Therefore, a LUN can be mapped to different device names after a Linux reboot or a reload of the external HBA drivers.

If you have more than one LUN, you need to have a persistent way of identifying each LUN. You can write a file system label to the LUN and mount the file system by that label.

**Viewing information
about LUNs**

To view information about LUNs, use the `sanlun` command.

The `sanlun` command is included in the FC Linux Host Utilities. The command displays the host device names and the storage system LUNs to which they are mapped.

Using the sanlun command: To view how host devices are mapped to storage system LUNs, complete the following step.

The `sanlun` program is installed from FC Linux Host Utilities.

Step	Action
1	Enter the following command at the host’s shell prompt: <pre># sanlun lun show all</pre> Result: The <code>sanlun</code> command displays information about FC devices and the storage system LUNs they are mapped to, including LUN size and state. Example: The following lines show sample <code>sanlun</code> command output: <pre># sanlun lun show all filer: lun-pathname device filename adapter protocol lun size lun state filer1: /vol/vol_1/lun_1 /dev/sdb1 host0 FCP 4g (4294967296) GOOD filer1: /vol/vol_1/lun_2 /dev/sdc1 host0 FCP 4g (4294967296) GOOD</pre>

Accessing LUNs without dm-multipath

Creating a partition

About partitions You can divide a LUN into one or more partitions. Each partition can contain a file system.

Partitions are required only if you use file system labels to maintain device persistency. If you want to use the entire LUN for a single file system, it is highly recommended to create a partition.

Before you begin You should have one or more LUNs available to the host. Log in to the host as root and run the following command to verify the LUNs are available:

sanlun lun show all

You should see a list of available LUNs. Verify that at least one LUN is listed.

Using fdisk to create a partition To create a partition of a LUN, complete the following steps.

Step	Action
1	Enter fdisk device . <i>device</i> is the device name assigned to the LUN, such as /dev/sda.
2	Enter the command n to create a new partition.
3	Enter p to create a primary partition.
4	Enter a partition number. For the first partition, enter 1 .
5	To accept the default cylinder values, press Enter twice.
6	Enter p to print the partition information. Disk /dev/sdb: 64 heads, 32 sectors , 10240 cylinders Units = cylinders of 2048 * 512 bytes Device Boot Start End Blocks Id System /dev/sdb1 1 10240 10485744 83 Linux

Step	Action	
7	If...	Then...
	The sectors value from Step 6 is not a multiple of 8	You must set the sectors manually to ensure best performance. Enter q to quit without creating the new partition. For the latest information and specific <code>fdisk</code> parameters, see bug number 156121 on Bugs Online at http://now.netapp.com/NOW/cgi-bin/bol.
	The sectors value from Step 6 is a multiple of 8	Enter w to save the changes and exit <code>fdisk</code> .

Creating a file system

About file systems

A file system enables users and applications to read and write files on a LUN. Linux supports several different file systems. You can create any type of file system. This section explains how to create a journaling ext3 file system.

Using mkfs to create a file system

To create a file system on the partition you created, run the `mkfs` command as shown and complete the following step.

Step	Action
1	Enter the following command: <pre>mkfs -t ext3 device</pre> <i>device</i> is the partition. For example, <code>/dev/sda1</code> is the first partition on <code>/dev/sda</code>.

Accessing LUNs without dm-multipath

Labeling and mounting the file system

About labels

A file system label creates a persistent mapping between a LUN and a file system. See “[Persistently identifying LUNs](#)” on page 66 for more information.

Once you create a label, you add it to the `/etc/fstab` file.

Using e2label to create a file system label

To create a file system label, run the `e2label` command and complete the following step.

Step	Action
1	Enter <code>e2label device label</code> <i>device</i> is the partition where you created the file system. <i>label</i> is a string (maximum 16 characters) used to identify the file system. Example: <code>/sbin/e2label /dev/sda1 lun0_label</code>

Mounting the file system

To mount the file system using the label you created, complete the following steps.

Step	Action
1	Create a mountpoint for the file system. Example: <code>mkdir /mnt/filer0_lun0</code>
2	Open the file <code>/etc/fstab</code> using a text editor.

Step	Action
3	Add an entry to /etc/fstab as follows, separating each entry with a tab (except the final two zeros, which are separated by a space): <pre><code>LABEL=<i>label</i> <i>mount_point</i> auto defaults 0 0</code></pre> <i>label</i> is the file system label you created. <i>mount_point</i> is the mountpoint you created for the file system. Example: <pre><code>LABEL=lun0_label /mnt/filer0_lun0 auto defaults 0 0</code></pre>
4	Save your changes to /etc/fstab and exit the editor.
5	Enter mount LABEL=<i>label</i> to mount the file system. The file system is automatically mounted whenever the Linux host is rebooted.
6	Do either of the following: ◆ Enter df to display the mounted file systems. Verify the new file system is listed. ◆ Change to the mountpoint directory and try writing and reading a file to verify the file system mounted successfully.

Unloading the external HBA driver and shutting down

Before unloading the external HBA driver

Before you unload the external HBA driver, ensure that the external HBA driver is not in use by unmounting any file systems on Fibre Channel LUNs, making offline the LVM volumes if any, flushing the multipath maps, stopping the dm-multipath daemon, and so on. You can verify whether external HBA driver is in use by using the `lsmod` command.

Similarly if you are using Veritas Storage Foundation, you need to stop all veritas processes using the driver and ensure that the driver is not in use by using the `lsmod` command.

Unloading the external HBA driver

To unload the external HBA driver, complete the following steps.

Step	Action
1	QLogic driver: Enter the appropriate command for your HBA model: <pre>modprobe -r qla2500 modprobe -r qla2400 modprobe -r qla2300</pre> Emulex driver: Enter the following command: <pre>modprobe -r lpfc</pre>

Shutting down the Linux host

Before shutting down a Linux host with Fibre Channel LUNs, flush the multipath maps by entering the following command:

```
multipath -F
```

Unloading the bundled HBA driver and shutting down

Before unloading the bundled HBA driver

Before you unload the bundled HBA driver, ensure that the bundled HBA driver is not in use by unmounting any file systems on Fibre Channel LUNs, making offline the LVM volumes if any, flushing the multipath maps, stopping the dm-multipath daemon, and so on. You can verify whether bundled HBA driver is in use by using the `lsmod` command.

Unloading the bundled HBA driver

To unload the bundled HBA driver, complete the following steps.

Step	Action
1	Bundled QLogic driver: Enter the appropriate command for your HBA model: If you are using Red Hat Enterprise Linux 4: <pre>modprobe -r qla2500 modprobe -r qla2400 modprobe -r qla2300</pre> If you are using Red Hat Enterprise Linux 5: <pre>modprobe -r qla2xxx</pre>
	Bundled Emulex driver: Enter the following command: <pre>modprobe -r lpfc</pre>

Shutting down the Linux host

Before shutting down a Linux host with Fibre Channel LUNs, flush the multipath maps by entering the following command:

```
multipath -F
```

About this chapter

This chapter applies only to SUSE Linux Enterprise Server hosts. For Red Hat Enterprise Linux hosts, see Chapter 6, “[Accessing LUNs from a Red Hat Enterprise Linux Host](#),” on page 51.

The steps for accessing LUNs depend on whether you are using dm-multipath or Veritas DMP support. This chapter has three main sections, one for accessing LUNs with dm-multipath, one for accessing LUNs with Veritas DMP, and one for accessing LUNs without dm-multipath.

Note

You must use dm-multipath or Veritas DMP if there is more than one physical path from the host to a LUN.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Loading the external HBA driver on the host](#)” on page 84
- ◆ “[Loading the bundled HBA driver on the host](#)” on page 85
- ◆ “[Configuring dm-multipath](#)” on page 86
- ◆ “[Starting multipath service](#)” on page 91
- ◆ “[Accessing LUNs with dm-multipath support](#)” on page 92
- ◆ “[Stopping dm-multipath](#)” on page 101
- ◆ “[Accessing LUNs with Veritas DMP support](#)” on page 102
- ◆ “[Accessing LUNs without dm-multipath support](#)” on page 103
- ◆ “[Unloading the external HBA driver and shutting down](#)” on page 111

Loading the external HBA driver on the host

Loading the external HBA driver

Once you have at least one LUN mapped to the Linux host, you can load the external HBA driver. When the driver is loaded, it scans the storage systems and discover all mapped LUNs. If you configure the dm-multipath or Veritas DMP support, multipath devices are created for all LUNs discovered.

Before you load the driver

Be sure that you have at least one LUN mapped to the Linux host as LUN 0 before loading the external HBA driver.

Starting FC

To start FC on a Linux host, load the driver for your HBAs.

Be sure you have configured multipathing before starting FC.

Loading the external QLogic HBA driver: To load the external QLogic HBA driver, complete the following step.

Step	Action
1	Enter the following command on the Linux host: <code>modprobe -v qla2300</code> <code>modprobe -v qla2400</code> <code>modprobe -v qla2500</code>

Loading the external Emulex HBA driver: To load the external Emulex HBA driver, complete the following step.

Step	Action
1	Enter the following command on the Linux host: <code>modprobe -v lpfc</code>

Loading the bundled HBA driver on the host

Loading the bundled HBA driver

Once you have at least one LUN mapped to the Linux host, you can load the bundled HBA driver. When the bundled HBA driver is started, it scans the storage systems and discovers all mapped LUNs. If you configure dm-multipath or Veritas DMP support, multipath devices are created for all LUNs discovered.

Before you load the bundled driver

Be sure that you have at least one LUN mapped to the Linux host as LUN 0 before loading the bundled HBA driver.

Starting FC

To start FC on a Linux host, load the bundled driver for your HBAs.
Be sure you have configured multipathing before starting FC.

Loading the bundled QLogic HBA driver: To load the bundled QLogic HBA driver, complete the following step.

Step	Action	
1	Enter the following command on the Linux host:	
	If...	Then...
	You are using SUSE Linux Enterprise Server 10	<code>modprobe -v qla2xxx</code>

Loading the bundled Emulex HBA driver: To load the bundled Emulex HBA driver, complete the following step.

Step	Action
1	Enter the following command on the Linux host: <code>modprobe -v lpfc</code>

Configuring dm-multipath

Note—
Skip this section if you are configuring LUNs without dm-multipath support (“[Accessing LUNs without dm-multipath support](#)” on page 103) or if you are using Veritas Storage Foundation (“[Accessing LUNs with Veritas DMP support](#)” on page 102).

Editing the configuration file

To configure dm-multipath, complete the following steps.

Step	Action	
1	Check if the /etc/multipath.conf file exists.	
	If...	Then...
	The file exists	a. Open the file. b. Comment out all the lines. c. Append the sample given “SUSE Linux Enterprise Server 10 SP2 and later:” on page 88 or “SUSE Linux Enterprise Server 9 SP3” on page 89, depending on your version of the operating system.
	The file does not exist	Create a new file with the entries as given in the sample “ SUSE Linux Enterprise Server 10 SP2 and later: ” on page 88 or “ SUSE Linux Enterprise Server 9 SP3 ” on page 89, depending on your version of the operating system.
2	Save the changes.	

Step	Action
3	Edit the <code>/etc/multipath.conf</code> file to exclude (blacklist) all devices that do not correspond to any LUN and which are mapped to your Linux host (that is, blacklist those devices that are not displayed in the sanlun lun show command). Complete the following steps to blacklist the devices. - Blacklist the non NetApp SCSI devices installed on your host, using the WWID method. To obtain the WWID of a device, run the scsi_id command on the respective device. Example: Suppose <code>/dev/sda</code> is a local SCSI drive. To obtain its WWID, run the following command: <pre>scsi_id -gus /block/sda</pre> The output looks like as given below: <pre>SIBM-ESXSMAW3073NC_FDAR9P66067WJ</pre> Add this entry for blacklisting. - Blacklist other devices using the devnode method by adding the following: <pre>devnode "^hd[a-z]"</pre> <pre>devnode "^(ram raw loop fd md dm- sr scd st)[0-9]*"</pre> <pre>devnode "^cciss!c[0-9]d[0-9]*"</pre> Example: On a SUSE Linux Enterprise Server 9 host: <pre>devnode_blacklist { devnode "^(ram raw loop fd md dm- sr scd st)[0-9]*" devnode "^hd[a-z][0-9]*" devnode "^cciss!c[0-9]d[0-9]*p[0-9]*" wwid SIBM-ESXSMAW3073NC_FDAR9P66067WJ }</pre>

Step	Action
4	If you are using ALUA, specify the ALUA callout program, that is, the prio "alua" program, instead of prio "ontap". Example: prio "alua" ALUA callout programs are available in the SUSE Linux Enterprise Server distribution.
5	Save the changes.

The following example shows the required changes to the /etc/multipath.conf file.

SUSE Linux Enterprise Server 10 SP2 and later:

```
defaults
{
    user_friendly_names    no
}
blacklist
{
    wwid <DevId>
    devnode  "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode  "^hd[a-z]"
    devnode  "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
}
devices
{
    device
    {
        vendor            "NETAPP"
        product           "LUN"
        getuid_callout     "/sbin/scsi_id -g -u -s /block/%n"
        prio              "ontap"
        features           "1 queue_if_no_path"
        hardware_handler   "0"
        path_grouping_policy group_by_prio
        failback           immediate
        rr_weight          uniform
        rr_min_io          128
        path_checker        directio
    }
}
```

<DevID> is a unique ID. See Step 3 of [“Editing the configuration file”](#) on page 86 to generate wwid.

SUSE Linux Enterprise Server 9 SP3:

```
defaults
{
    user_friendly_names    no
}
devnode_blacklist
{
    wwid        <DevID>
    devnode    "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode    "^hd[a-z]"
    devnode    "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
}
devices
{
    device
    {
        vendor            "NETAPP"
        product           "LUN"
        getuid_callout    "/sbin/scsi_id -g -u -s /block/%n"
        prio_callout     "/sbin/mpath_prio_ontap /dev/%n"
        features          "1 queue_if_no_path"
        hardware_handler  "0"
        path_grouping_policy group_by_prio
        failback          immediate
        rr_weight         uniform
        rr_min_io         128
        path_checker      readsector0
    }
}
```

For more information about tuning the multipath configuration, see [“Tuning the dm-multipath configuration”](#) on page 100.

Default versus recommended settings

The default versus recommended settings for SUSE Linux Enterprise Server 9 and SUSE Linux Enterprise Server 10 are as follows:

Component	Parameter	SUSE Linux Enterprise Server 9		SUSE Linux Enterprise Server 10	
		Default value	NetApp recommended value	Default value	NetApp recommended value
QLogic External Driver	MaxRetriesPerIo	0	10	0	10
	MaxRetriesPerPath	0	10	0	10
	ql2xsuspendcount	10	6	10	6
QLogic Bundled Driver	qlport_down_retry	0	30	0	30
	ql2xsuspendcount	10	6	NA	NA
	MaxRetriesPerIo	0	10	NA	NA
	MaxRetriesPerPath	0	10	NA	NA
QLogic FC HBA	Link Down Timeout	30	20	30	20
Multipath.conf	user_friendly_names	no	no	no	no
Veritas DMP	dmp_restore_interval	300	60	300	60

Starting multipath service

Note

Skip this section if you are configuring LUNs without dm-multipath support ([“Accessing LUNs without dm-multipath support”](#) on page 103) or if you are using Veritas Storage Foundation ([“Accessing LUNs with Veritas DMP support”](#) on page 102).

Once you have loaded the external HBA drivers, start the multipath service on the Linux host. Complete the following steps.

Step	Action
1	To start the multipath service, enter the following command: <code>/etc/init.d/boot.multipath start</code>
2	To start the multipath daemon, enter the following command: <code>/etc/init.d/multipathd start</code>

Configuring the multipath service to start automatically after

reboot: Once you have started the multipath service on the Linux host, you can configure it to start automatically after reboot. Complete the following step.

Step	Action
1	Add the multipath service to the boot sequence by entering the following commands: <code>insserv boot.multipath multipathd</code> <code>chkconfig --add multipathd</code> <code>chkconfig multipathd on</code>

Accessing LUNs with dm-multipath support

Note

Skip this section if you are configuring LUNs without dm-multipath support ([“Accessing LUNs without dm-multipath support”](#) on page 103) or if you are using Veritas Storage Foundation ([“Accessing LUNs with Veritas DMP support”](#) on page 102).

To access LUNs using dm-multipath support, perform the following steps.

1. Verify the HBA setting and dm-multipath configuration. For more information, see [“Verifying the HBA setting and dm-multipath configuration”](#) on page 93.
2. Identify and view the list of LUNs. For more information, see [“Identifying and viewing a list of LUNs”](#) on page 96.
3. Create and mount the file system. For more information, see [“Creating and mounting a file system on a LUN”](#) on page 98.
4. Tune the dm-multipath configuration if required. For more information, see [“Tuning the dm-multipath configuration”](#) on page 100.

Verifying the HBA setting and dm-multipath configuration

Overview

Once you start dm-multipath, you should verify that the desired multipathing policy is correctly configured, and the multipath maps were created.

Verifying the HBA setting

QLogic:

Verify the ql2xsuspendcount, MaxRetriesPerIo, and MaxRetriesPerPath parameters in the following way:

Check the values for these parameters in the corresponding qla2xxx directory in the sysfs, /sys/module/qla2xxx/ or /sys/module/qla2xxx/parameters/.

Example: You can check the value of ql2xsuspendcount on a SUSE Linux Enterprise Server host by running the following command:

```
cat /sys/module/qla2xxx/ql2xsuspendcount
```

Note

You cannot verify the same in SUSE Linux Enterprise Server 9 as there is no qla2xxx directory available in sysfs.

PortDownRetryCount and LinkDownTimeout along with Connection Options and Data Rate (which need to manually set instead of the default autonegotiate) can be verified as follows.

Run the **-c all** option of the **scli** command:

```
# scli -c all
```

Example: The following is a sample output of the **scli** command:

```
-----
HBA 0: QLE2462 Port 1 WWPN 21-00-00-E0-8B-80-72-A1 Port ID 03-05-00
-----
Connection Options : 1 - Point-to-Point Only.
Data Rate : 4 Gbps
Frame Size : 2048
Hard Loop ID : 0
Loop Reset Delay (seconds) : 5
Enable Host HBA BIOS : Disabled
Enable Hard Loop ID : Disabled
Enable FC Tape Support : Enabled
```

```

Operation Mode : 0 - Interrupt for every I/O completion
Interrupt Delay Timer (100ms) : 0
Execution Throttle : 16
Login Retry Count : 8
Port Down Retry Count : 30
Enable LIP Full Login : Enabled
Link Down Timeout (seconds) : 20
Enable Target Reset : Enabled
LUNs Per Target : 128
-----
HBA 1: QLE2462 Port 2 WWPN 21-01-00-E0-8B-A0-72-A1 Port ID 03-01-00
-----
Connection Options : 1 - Point-to-Point Only.
Data Rate : 4 Gbps
Frame Size : 2048
Hard Loop ID : 0
Loop Reset Delay (seconds) : 5
Enable Host HBA BIOS : Disabled
Enable Hard Loop ID : Disabled
Enable FC Tape Support : Enabled
Operation Mode : 0 - Interrupt for every I/O completion
Interrupt Delay Timer (100ms) : 0
Execution Throttle : 16
Login Retry Count : 8
Port Down Retry Count : 30
Enable LIP Full Login : Enabled
Link Down Timeout (seconds): 20
Enable Target Reset : Enabled
LUNs Per Target : 128

```

Emulex:

The default settings work fine in case of Emulex.

Verifying the configuration

Use the **multipath** command on the Linux host to view the dm-multipath configuration. You can change the amount of detail displayed by using the **-v** option. For details on the multipath command options, enter **multipath -h**.

To view all of the configuration details, enter the following command on the Linux host:

```
multipath -v3 -d -ll
```

Note

The `-d` (“dry run”) parameter prevents the command from actually updating the multipath maps.

Verify that the `multipathd` daemon is running by issuing the following command:

```
/etc/init.d/multipathd status
```

Viewing the multipath devices: To view a list of multipath devices, including which `/dev/sdx` devices are used, enter the following command:

```
multipath -ll
```

Example: The following output shows two LUNs, with four paths each—two proxy and two non-proxy paths.

```
360a98000486e5372635a44646a505643NETAPP,LUN
[size=10G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=8][active]
  \_ 6:0:0:0 sda 8:0 [active][ready]
  \_ 7:0:1:0 sdg 8:96 [active][ready]
\_ round-robin 0 [prio=2][enabled]
  \_ 6:0:1:0 sdc 8:32 [active][ready]
  \_ 7:0:0:0 sde 8:64 [active][ready]
360a98000486e5363693444646a2f656cNETAPP,LUN
[size=10G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=8][enabled]
  \_ 6:0:1:1 sdd 8:48 [active][ready]
  \_ 7:0:0:1 sdf 8:80 [active][ready]
\_ round-robin 0 [prio=2][enabled]
  \_ 6:0:0:1 sdb 8:16 [active][ready]
  \_ 7:0:1:1 sdh 8:112 [active][ready]
```

Note

The `/dev/mapper` devices are persistent across reboots, but the `/dev/sdx` devices are not. After a reboot, or a reload of the external HBA driver, you might find that different `/dev/sdx` devices make up a specific `/dev/mapper` device. The `/dev/mapper` device, however, will always correspond to the same LUN.

To view the `/dev/mapper` devices, enter the following command:

```
ls -l /dev/mapper/
```

Example: The following is a sample output.

```
brw----- 1 root root 253, 1 Sep 20 17:09
360a98000486e5363693444646a2f656c
brw----- 1 root root 253, 0 Sep 20 17:09
360a98000486e5372635a44646a505643
```

Identifying and viewing a list of LUNs

Persistently identifying LUNs

To create partition, you need to identify that particular LUN.

The dm-multipath code creates persistent devices for each LUN in the `/dev/mapper/` directory on the Linux host. Use these entries for multipath LUNs. You can create a file system directly on a multipath device in `/dev/mapper/`. There is no need to create a partition or label on the multipath device.

The `/dev/mapper/` entries can also be used in the `/etc/fstab` file.

Viewing information about LUNs

To view information about LUNs, use the **sanlun** command.

Using sanlun: To view a list of LUNs that are mapped to the Linux host, enter the following command on the Linux host console:

```
sanlun lun show all
```

Example: The following output shows one LUN with four paths mapped to the Linux host.

```
# sanlun lun show all
filer: lun-pathname device filename adapter protocol lun size lun
state
f3020-92: /vol/vol24/lun24 /dev/sdt host4 FCP 5g (5368709120) GOOD
f3020-93: /vol/vol24/lun24 /dev/sdu host4 FCP 5g (5368709120) GOOD
f3020-93: /vol/vol24/lun24 /dev/sdk host3 FCP 5g (5368709120) GOOD
f3020-92: /vol/vol24/lun24 /dev/sdj host3 FCP 5g (5368709120) GOOD
```

To display the multipath information, for native device-mapper multipath, enter the following command:

```
sanlun lun show -p
```

```
# sanlun lun show -p
f3020-92:/vol/vol1/lun2 (LUN 1)          Lun state: GOOD
Lun Size:      3.6g (3879731200)    Filer_CF_State: Cluster Enabled
Protocol: FCP          Filer Partner: f3020-93
DM-MP DevName: mpath3    (360a9800043346852563444717a513571) dm-1
Multipath-provider: NATIVE
```

sanlun	Filer			Primary	Partner
path	Path	/dev/	Host	Filer	Filer
state	type	node	HBA	port	port
UP	primary	sde	host17	1b	--
UP	primary	sdg	host18	1a	--
UP	secondary	sdh	host17	--	1a
UP	secondary	sdi	host18	--	1b

Accessing LUNs with dm-multipath support

Creating and mounting a file system on a LUN

Overview

Once a LUN is added to the multipath configuration, you can create a file system.

When creating a file system, you use the multipath device in the /dev/mapper/ directory.

Creating the file system

To create a file system on a LUN, enter the following command on the Linux host console:

mkfs -t *type* /dev/mapper/*device*

type is the file system type, such as ext3 .

device is the multipath device name of the LUN in the /dev/mapper/ directory.

Example: The following command creates an ext3 file system on a LUN:

mkfs -t ext3 /dev/mapper/360a9800043346536684a337a7553424f

Mounting the file system

Multipath devices in the /dev/mapper/ directory that are listed in /etc/fstab are automatically remounted after a reboot, assuming you have correctly configured the multipath service to start automatically.

To mount the new file system, you create a mountpoint, and add an entry to the /etc/fstab file. To mount the file system, complete the following steps.

Step	Action
1	Create a mountpoint using the following command: mkdir <i>mount_point</i> <i>mount_point</i> is the name of the mountpoint you are creating.
2	Open the /etc/fstab file with a text editor.

Step	Action
3	Add an entry to <code>/etc/fstab</code> as follows, separating each entry with a tab (except the final two zeros, which are separated by a space): <pre>device mount_point type defaults 0 0</pre> <i>device</i> is the name in the <code>/dev/mapper/</code> directory. <i>mount_point</i> is the mountpoint you created for the file system. <i>type</i> is the file system type, such as <code>ext2</code>, <code>ext3</code>, or <code>reiserfs</code>. Example: <pre>/dev/mapper/360a9800043346536684a337a7553424f /mnt/ssl/lun_0 ext3 defaults 0 0</pre>
4	Enter the following command to mount the file system: <pre>mount -a</pre>
5	Do either of the following: ◆ Enter df to display the mounted file systems. Verify the new file system is listed. ◆ Change to the mountpoint and try writing and reading a file to verify the file system mounted successfully.

Tuning the dm-multipath configuration

Tuning multipath parameters

Advanced administrators who are familiar with Linux performance tuning may want to modify one of the Linux parameters for best performance.

queue_depth: The queue_depth value for the SCSI device queue is one of the critical settings to obtain the best performance. This value can be found in sysfs, and you can echo a new value in to change this for each device SCSI command queue. The following example shows how to change the SCSI device queue depth for device sdb.

```
# cat /sys/block/sdb/device/queue_depth
32
# echo "64" > /sys/block/sdb/device/queue_depth
# cat /sys/block/sdb/device/queue_depth
64
```

rr_min_io: The rr_min_io setting for dm-multipath specifies the number of I/O operations sent through a path before switching to the next path. Lowering this value from the default value of 1000 has been seen to dramatically affect overall throughput for dm-multipath, especially for large I/O workloads (64 KB or more). Other performance tests have shown a good value for rr_min_io to be N, where N is less than or equal to the number of threads in the system that issue a single I/O and block, waiting for it to complete. Tuning the rr_min_io value is suggested for optimal performance because it really depends on the workload and other factors such as the host memory and CPU speed.

The rr_min_io can be set in the device section of the multipath.conf file as follows:

```
device
{
    rr_min_io      128
}
```

See the “[Configuring dm-multipath](#)” on page 86 for details on blacklisting the local hard drives and other resources that should not be included in the multipathing configuration.

Stopping dm-multipath

Note

Skip this section if you are configuring LUNs without dm-multipath support (“[Accessing LUNs without dm-multipath support](#)” on page 103) or if you are using Veritas Storage Foundation (“[Accessing LUNs with Veritas DMP support](#)” on page 102).

When you want to stop the dm-multipath support on the Linux host, you should stop the affected services in the order shown in this section.

To stop dm-multipath and FC on the host, complete the following steps.

Step	Action
1	Unmount all file systems on the FC LUNs.
2	Flush the multipath maps by entering the following command: multipath -F
3	Stop the multipath daemon by entering the following command: /etc/init.d/multipathd stop
4	Unload all of the dm-multipath modules by entering the following command: /etc/init.d/boot.multipath stop

Accessing LUNs with Veritas DMP support

Veritas Storage Foundation

If you are using Veritas Storage Foundation, use the Veritas DMP solution. For details about configuring the Veritas Volume Manager and accessing LUNs, see the *Veritas® Volume Manager Administrator's Guide* for Linux that is shipped along with the software.

The `vxddmp restore` daemon should be configured with an interval of 60.

To configure the daemon, run the following command:

```
vxddmptune dmp_restore_interval 60
```

On reboot this value comes into effect and remains persistent.

You can also use the following commands to configure the daemon value.

```
vxddmpadm stop restore
```

```
vxddmpadm start restore interval=60
```

This value comes into effect immediately and is not persistent on reboot.

To verify that the daemon value has been successfully configured, run the following command:

```
vxddmpadm stat restored
```

Accessing LUNs without dm-multipath support

Overview

Without multipathing software, the host has no way of distinguishing multiple LUNs from multiple paths to the same LUN.

If you do not use dm-multipath or Veritas DMP, you must limit each LUN to a single path. If you are using an active/active configuration storage system, you must use the standby cfmode to avoid having a secondary path visible on the storage system.

If you are using a switched Fibre Channel fabric, you must zone your switch to prevent the host from having multiple paths to any one LUN. See “[Zoning with standby or mixed cfmode](#)” on page 48. You can create a switch zone with a single host HBA WWPN and a single storage system WWPN; this zone provides a single path to a LUN.

There are different HBA parameter settings for configurations without dm-multipath. See “[Setting the required HBA and driver parameters](#)” on page 24.

Configurations without dm-multipath can have additional restrictions on host OS versions and HBA models. See the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Tasks for accessing LUNs

After the initiator itself is configured, you need to set up the host to access the LUNs you created on the storage system. To access LUNs, you must complete the following tasks.

1. Identify and view the list of LUNs. For more information, see “[Identifying and viewing the LUN information](#)” on page 105.
2. Create one or more partitions of the LUN. A file system needs a partition, even if the one partition uses all of the space in the LUN. For more information, see “[Creating a partition](#)” on page 107.
3. Create a file system on the partition. For more information, see “[Creating a file system](#)” on page 108.

4. Create a label for the file system and add an entry for the file system's label in `/etc/fstab`. Mount the file system. For more information, see [“Labeling and mounting the file system”](#) on page 109.

Identifying and viewing the LUN information

Persistently identifying LUNs

LUNs on the storage system are seen by the Linux host as SCSI devices. Linux assigns a unique device name, such as `/dev/sda` or `/dev/sdb`, to each LUN. The device names are assigned as the LUNs are discovered. Because of variable network delays and other issues, the order in which LUNs are discovered cannot be predicted. Therefore, a LUN can be mapped to different device name after a Linux reboot or a reload of the external HBA drivers.

If you have more than one LUN, you need to have a persistent way of identifying each LUN. You can write a file system label to the LUN and mount the file system by that label.

Viewing information about LUNs

You can view information about LUNs in the following ways:

- ◆ `lsscsi`: The **`lsscsi`** command lists all SCSI devices including iSCSI devices and their attributes.
- ◆ `sanlun`: The **`sanlun`** command is included in FC Linux Host Utilities. It displays the host device names and the storage system LUNs to which they are mapped.

Using the `lsscsi` command

The **`lsscsi`** command lists information about all the SCSI devices and their attributes.

Example: The following example shows the output of the **`lsscsi`** command.

```
[0:0:0:0] cd/dvd HL-DT-ST RW/DVD GCC-4246N 0Y03 /dev/sr0
[2:0:0:0] disk IBM-ESXS PYH073C3-ETS10FN RXQN /dev/sda
[2:0:8:0] process IBM 25P3495a S320 1 1 -
[4:0:0:0] disk NETAPP LUN 0.2 /dev/sdd
[4:0:0:1] disk NETAPP LUN 0.2 /dev/sde
[5:0:0:0] disk NETAPP LUN 0.2 /dev/sdf
[5:0:0:1] disk NETAPP LUN 0.2 /dev/sdg
[6:0:0:0] disk NETAPP LUN 0.2 /dev/sdh
[6:0:0:1] disk NETAPP LUN 0.2 /dev/sdi
```

Using the sanlun command

To view how host devices are mapped to storage system LUNs, complete the following steps.

Step	Action
1	On the host, go to the /opt/netapp/santools directory.
2	Enter the following command: sanlun lun show all Result: The sanlun command displays information about FC devices and the storage system LUNs they are mapped to, including LUN size and state. Example: The following lines show sample sanlun command output. # sanlun lun show all filer: lun-pathname device filename adapter protocol lun size lun state f3020-92: /vol/vol24/lun24 /dev/sdt host4 FCP 5g (5368709120) GOOD f3020-93: /vol/vol24/lun24 /dev/sdu host4 FCP 5g (5368709120) GOOD f3020-93: /vol/vol24/lun24 /dev/sdk host3 FCP 5g (5368709120) GOOD f3020-92: /vol/vol24/lun24 /dev/sdj host3 FCP 5g (5368709120) GOOD

Creating a partition

About partitions You can divide a LUN into one or more partitions. Each partition can contain a file system.

Partitions are required only if you use file system labels to maintain device persistency. If you want to use the entire LUN for a single file system, it is highly recommended to create a partition.

Before you begin You should have one or more LUNs available to the host. Log in to the host as root and run the following command to verify the LUNs are available:

sanlun lun show all

You should see a list of available LUNs. Verify that at least one LUN is listed.

Using fdisk to create a partition To create a partition of a LUN using fdisk, complete the following steps.

Step	Action
1	Enter the following command: fdisk fcp_device <i>fcp_device</i> is the link that maps the FC LUN on the storage system to a local disk device.
2	Enter the command n to create a new partition.
3	Enter p to create a primary partition.
4	Enter a partition number. For the first partition, enter 1 .
5	Press Enter twice to accept the default cylinder values.
6	Enter w to save the changes and exit fdisk.

Creating a file system

About file systems A file system enables users and applications to read and write files on a LUN. Linux supports several different file systems. You can create any type of file system. This section explains how to create a journaling ext3 file system.

Creating a file system To create a file system on a LUN, complete the following step.

Step	Action
1	On the host console, enter the following command: <pre>mkfs -V -t <i>filesystem_type</i> <i>partition</i></pre> <i>partition</i> is the name of the partition you just created. For SUSE Linux Enterprise Server 9, the partition name is the name of the link that maps the LUN on the storage system to a local disk device concatenated with the letter p and the partition number. <i>filesystem_type</i> is the type of file system you want to create. Example: The following example creates an ext3 file system on the first partition of LUN 0. <pre>mkfs -V -t ext3 /dev/disk/by-id/360a98000433464624f4a34732f785a47-0p1</pre>

Accessing LUNs without dm-multipath

Labeling and mounting the file system

About labels

A file system label creates a persistent mapping between a LUN and a file system. See “[Persistently identifying LUNs](#)” on page 105 for more information.

Once you create a label, you add it to the /etc/fstab file.

Using e2label to create a file system label

To create a file system label, complete the following step.

Step	Action
1	Enter e2label device label <i>device</i> is the partition where you created the file system. <i>label</i> is a string (maximum 16 characters) used to identify the file system. Example: <code>/sbin/e2label /dev/sda1 lun0_label</code>

Mounting the file system

To mount the file system using the label created, complete the following steps.

Step	Action
1	On the host, create a mountpoint for the LUNs.
2	With a text editor, open the host’s /etc/fstab file.
3	Add an entry for the file system you created. Example: <code>/dev/disk/by-id/360a98000433464624f4a34732f785a47-0p1 /mnt_pt ext3 defaults 0 0</code> <i>mnt_pt</i> is the mountpoint you created in Step 1. Result: The file systems listed are mounted when Linux starts.
4	Save your changes and exit the /etc/fstab file.

Step	Action
5	To mount the file systems, enter the following command on the host console: mount

Unloading the external HBA driver and shutting down

Before unloading the external HBA driver

Before you unload the external HBA driver, ensure that the external HBA driver is not in use by unmounting any file systems on Fibre Channel LUNs, making offline the LVM volumes if any, flushing the multipath maps, stopping the dm-multipath daemon, and so on. You can verify whether external HBA driver is in use by using the `lsmod` command.

Similarly if you are using Veritas Storage Foundation, you need to stop all veritas processes using the driver and ensure that the driver is not in use by using the `lsmod` command.

Unloading the external HBA driver

To unload the external HBA driver, complete the following steps.

Step	Action
1	QLogic driver: Enter the appropriate command for your HBA model: <pre>modprobe -r qla2500 modprobe -r qla2400 modprobe -r qla2300</pre> Emulex driver: Enter the following command: <pre>modprobe -r lpfc</pre>

Shutting down the Linux host

Before shutting down a Linux host with Fibre Channel LUNs, flush the multipath maps by entering the following command:

```
multipath -F
```

Note

For information about the supported configuration, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

About this chapter

SAN boot uses a SAN-attached disk, such as a LUN, as a boot device for a host. This chapter provides the tasks for configuring SAN boot.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Configuring SAN boot on the native host](#)” on page 114
- ◆ “[Configuring SAN boot for Veritas Storage Foundation](#)” on page 121

Configuring SAN boot on the native host

Tasks for configuring SAN boot

To configure SAN boot, perform the following tasks.

- 1. Install the operating system on the LUN configured on the storage controller.
- 2. Install the Host Utilities.
- 3. Install and configure the external HBA driver or the bundled HBA driver.
- 4. Configure dm-multipath according to the recommended procedure.

Installing the operating system

To install the operating system on SAN boot LUN, complete the following steps.

Step	Action
1	On the storage system, choose a cfmode setting.
2	Create a LUN on the storage system which acts as the SAN boot LUN and map it to the host. Attention _____ Ensure that the SAN boot LUN is mapped as LUN ID 0 for the corresponding host, and multiple paths to the SAN boot LUN are available on the host. Note _____ Ensure the SAN boot LUN is visible to the host during booting up.
3	Enable the BIOS of the HBA port to which the SAN boot LUN is mapped. For information about how to enable the HBA BIOS, see your HBA vendor-specific documentation. Note _____ It is best to enable the spinup delay option for the individual HBA ports in the HBA BIOS.

Step	Action
4	Configure the paths to the HBA boot BIOS as primary, secondary, tertiary, and so on, on the boot device. For more information, see your vendor-specific documentation.
5	After making the relevant changes to the HBA BIOS, save and exit.
6	Reboot the host.
7	Install the operating system on the LUN configured on storage controller. Attention _____ Specify <i>linux mpath</i> in the anaconda installer, during the operating system installation. When <i>linux mpath</i> is specified, you can see the multipath devices (/dev/mapper/mpath) as installation devices. _____
8	After the operating system installation, reboot the host.

Installing Host Utilities

Install FC Linux Host Utilities 4.2. For information about installing the FC Linux Host Utilities 4.2, see “[Installing the FC Linux Host Utilities software](#)” on page 16.

Installing and configuring the external HBA driver

To install and configure the external HBA driver on the host, complete the following steps.

Step	Action
1	Download and install the HBA driver and applications. For more information, see Appendix , “Downloading and installing the external HBA driver and applications,” on page 22. Attention _____ Run the qlinstall script with the -up option to install the external QLogic driver on the SAN booted host. Reboot the host for loading the above external driver. _____

Step	Action
2	Set the required HBA and driver parameters. QLogic: For more information about setting the QLogic HBA parameters, see Appendix , “QLogic HBA parameters,” on page 24. To set the HBA driver parameters, perform the following steps. - Edit the <code>/etc/modprobe.conf</code>. - Append the options <code>qla2xxx</code> statement with settings in the <code>modprobe.conf</code> file on a single line; only one options <code>qla2xxx</code> statement is allowed. <pre>ql2xsuspendcount=6 MaxRetriesPerPath=10 MaxRetriesPerIo=10 ql2xfailover=0</pre> - Save the file and exit. - Enter the appropriate command for your HBA model to rebuild the initial ramdisk (<code>initrd</code>) image with the new parameter values: <pre>./qlinstall -br -in qla2500 ./qlinstall -br -in qla2400 ./qlinstall -br -in qla2300</pre> - Reboot the host for loading the above external driver. Emulex: Attention _____ Run the <code>lpfc-install</code> script to install the external Emulex driver on the SAN booted host. _____ No changes are required.

Installing and configuring the bundled HBA driver

To install and configure the bundled HBA driver, complete the following steps.

Step	Action
1	Download and install the HBA vendor’s command-line interface (CLI) and graphical user interface (GUI) management packages from the respective HBA vendor Web site. Note These packages contains s/w components used to display HBA information (also used by the sanlun utility). For QLogic: Download and install the following from the QLogic Web site: ❖ SANsurfer FC HBA Manager (for sanlun fcp show adapter) ❖ SANsurfer CLI For Emulex: Download and install the following from the Emulex Web site: ❖ Applications kit: Ensure that the application kit version matches the bundled Emulex driver version.

Step	Action
2	QLogic: Configure the bundled HBA parameter. To configure the bundled HBA parameters, complete the following steps. - Run the following command: <code>scli</code> - Set the following HBA parameters: ❖ Port Down Retry Count to 30 ❖ Link Down Timeout to 20 - Configure the QLogic HBA port topology to Point to Point Only while connecting to a storage system operating in SSI, Standby or Partner CF modes, or to Loop Only if the storage system is operating in Dual Fabric or Mixed CF modes. - Configure the QLogic port speed ideally to the highest speed possible, depending on its ceiling and that of the switch or target port it is connected to. To do this select the Data Rate option and set it to the specified speed. - Select Commit Changes for this HBA so that the changes are applied to the HBA ports. - Repeat the settings for each HBA port listed.
3	For QLogic: No changes are required for the QLogic bundled driver in the modprobe.conf file, the default settings works fine. Note_____ There are no qla2300/qla2400/qla2500 firmware modules in the Red Hat Enterprise Linux 5 series. These modules have been integrated into the base qla2xxx module itself. For Emulex: No special settings are required for supported bundled Emulex HBAs. The default settings work fine in the case of Emulex.

Configuring dm-multipath

Note

You have to download the respective errata multipathing package from the Red Hat Network Web site at rhn.redhat.com/ for SAN boot support on Red Hat Enterprise Linux 5 Update 1. For more information about the errata package, see the Notes section of the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>.

Configure dm-multipath. For more information about configuring the dm-multipath, see “[Configuring dm-multipath](#)” on page 55.

To activate multipathing during booting, enter the following command on the SAN booted host:

```
chkconfig multipathd on
```

Configuring SAN boot for Veritas Storage Foundation

Configuring SAN boot on the host

To configure SAN boot on the host using a storage system, complete the following steps.

Step	Action
1	On the storage system, choose a supported cfmode.
2	Create a LUN on the storage system which will act as the SAN boot LUN.
3	Map the LUN to the host.
4	Ensure that only one non-proxy path is available to the LUN.
5	Ensure that only one SAN boot LUN is available to the host and that LUN is mapped as LUN ID 0.
6	Enable the boot BIOS of the HBA port to which the SAN boot LUN is mapped. For information about how to enable the boot BIOS, see the respective HBA vendor-specific documentation. Note _____ Enable the spinup delay option for the HBA port. This permits the HBA BIOS to scan for three minutes until it finds a boot device. _____
7	Note _____ Before installing the OS see the section on rootability in the <i>Veritas™ Volume Manager Administrator's Guide</i> for Linux that is shipped along with the software for partitioning information. _____ Attention _____ Red Hat Enterprise Linux 4 Update 4 does not have bundled HBA drivers for 4-Gb and 8-Gb QLogic cards, therefore use the device driver kit provided by QLogic. See NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability on supported drivers. _____ After performing the appropriate changes to the HBA BIOS and ensuring the SAN boot LUN is visible, install the operating system on the SAN boot LUN.

Step	Action
8	After the operating system installation, reboot the host. As a result, the host will boot from the SAN boot LUN on the storage system through a non-proxy path.
9	Install FC Linux Host Utilities 4.2. For information about installing the FC Linux Host Utilities 4.2, see “ Installing the FC Linux Host Utilities software ” on page 16.
10	Install the supported version of the drivers, if you are using external HBA drivers. For more information about the installation and configuration of the drivers, see Chapter 2, “ Installing HBAs and Software on the Host ,” on page 15.
11	Verify the HBA settings. For more information, see “ Verifying the HBA settings ” on page 63.
12	Install Veritas Storage Foundation, any relevant patches or fixes, and respective ASL and APM.
13	The vxddmp restore daemon should be configured with an interval of 60. To configure the daemon, run the following command: vxddmptune dmp_restore_interval 60 On reboot this value comes into effect and remains persistent.
14	Encapsulate the root disk for use in VxVM. Any existing data on the disk is preserved by encapsulation when the disk is placed under VxVM control. To encapsulate the disk run the following command: vxdiskadm For detailed steps, see the section on encapsulating the disk in <i>Veritas™ Volume Manager Administrator's Guide</i> for Linux that is shipped along with the software.
15	Reboot the host after encapsulation.
16	Verify the encapsulation by issuing the following command: vxprint It will display the rootvol and swapvol under the corresponding disk group.

Step	Action
17	Configure the paths to the HBA boot BIOS as primary, secondary, tertiary, and so on, on the boot device. For more information, see the respective HBA vendor-specific documentation. Note_____ It is recommended to have non-proxy paths as primary paths in the HBA BIOS. _____

Note

SAN boot is supported with SUSE Linux Enterprise Server and Veritas Storage Foundation. For information about the supported configuration, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

About this chapter

SAN boot uses a SAN-attached disk, such as a LUN configured on storage controller, as a boot device for a host. This chapter provides the tasks for configuring SAN boot.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Configuring SAN boot on the native host](#)” on page 126
- ◆ “[Configuring SAN boot for Veritas Storage Foundation](#)” on page 131

Configuring SAN boot on the native host

Overview of tasks for configuring SAN boot

The following is an overview of the tasks you need to perform to configure SAN boot with multipath support:

- 1. Install the operating system on the LUN configured on storage controller with root that is on a single SCSI device.
- 2. Configure the external or bundled HBA driver.
- 3. Configure root with dm-multipath support according to the recommended procedure.

Install the OS on the LUN configured on a storage controller with root on a single SCSI device

To install the operating system on the LUN configured on storage controller with root on a single SCSI device, complete the following steps.

Step	Action
1.	Create a LUN on the storage system which acts as the SAN boot LUN and map it to the host. Attention — Ensure that the SAN boot LUN is mapped as LUN ID 0 for the corresponding host, and only a single path to the SAN boot LUN is available on the host. Note — Ensure the SAN boot LUN is visible to the host during bootup.
2.	Enable the BIOS of the HBA port to which the SAN boot LUN is mapped. For information about how to enable the HBA BIOS, see your HBA vendor-specific documentation.
3.	After making the relevant changes to the HBA BIOS and ensuring the SAN booted LUN is visible, install the operating system on the SAN boot LUN.

Installing and configuring external HBA driver

To install and configure the external HBA driver on the host, complete the following steps.

Step	Action
1	Download and install the external HBA driver and applications. For more information, see Appendix , “Downloading and installing the external HBA driver and applications,” on page 22. Attention Run the qlinstall script with the -up option to install the external QLogic driver on the SAN booted host. Reboot the host for loading the earlier external driver.
2	Set the required HBA and driver parameters. QLogic: For more information about setting the QLogic HBA parameters, see “QLogic HBA parameters” on page 24. To set the external HBA driver parameters, perform the following steps. - Edit <code>/etc/modprobe.conf</code>. - Append the options <code>qla2xxx</code> statement with settings in the <code>modprobe.conf</code> file on a single line; only the option <code>qla2xxx</code> statement is allowed. <pre>ql2xsuspendcount=6 MaxRetriesPerPath=10 MaxRetriesPerIo=10 ql2xfailover=0</pre> - Save the file and exit. Emulex: Attention Run the lpfc-install script to install the external Emulex driver on the SAN booted host. No changes are required.

**Configuring
bundled HBA driver**

To configure the bundled HBA driver, complete the following steps.

Step	Action
1	Download and install the HBA vendor’s command-line interface (CLI) and graphical user interface (GUI) management packages from the respective HBA vendor Web site. Note These packages contain software components that are used to display HBA information (also used by the sanlun utility). For QLogic: Download and install the following from the QLogic Web site: ❖ SANsurfer FC HBA Manager (for sanlun fcp show adapter) ❖ SANsurfer CLI For Emulex: Download and install the following from the Emulex Web site: ❖ Applications kit: Ensure that the application kit version matches the bundled Emulex driver version.

Step	Action
2	Set the required bundled HBA parameters. QLogic: To configure the bundled HBA parameter, perform the following steps: - Run the following command: scli - Set the following HBA parameters: ❖ Port Down Retry Count to 30 ❖ Link Down Timeout to 20 - Configure the QLogic HBA port topology to Point to Point Only while connecting to a storage controller operating in SSI, Standby or Partner CF modes, or to Loop Only if the storage system is operating in Dual Fabric or Mixed CF modes. - Configure the QLogic port speed ideally to the highest speed possible, depending on its ceiling and that of the switch or target port it is connected to. To do this select the Data Rate option and set it to the specified speed. - Select Commit Changes for this HBA so that the changes are applied to the HBA ports. - Repeat the settings for each HBA port listed.
3	For QLogic: To set the bundled HBA driver parameters, perform the following steps: - Edit the modprobe.conf.local file. - Append the following line: options qla2xxx qlport_down_retry=30 - Save the file and exit. For Emulex: No changes are required.

Configure root with dm-multipath support

To configure root with dm-multipath, complete the following steps.

Step	Action
1	Configure dm-multipath. For more information about configuring dm-multipath, see “ Configuring dm-multipath ” on page 86.
2	Install the Linux FC Linux Host Utilities 4.2. For information about installing the FC Linux Host Utilities 4.2, see “ Installing the FC Linux Host Utilities software ” on page 16.
3	On the SAN booted host, run the following commands to activate multipathing during bootup: chkconfig boot.multipath on chkconfig multipathd on
4	Re-create the initrd by running the following command: mkinitrd -f mpath Attention _____ Keep a backup of the existing initrd image and make necessary modifications in the bootloader. _____
5	Enable all paths to the LUN on the storage controller.
6	Reboot the host. Note _____ As all paths to the bootlun are now available on the host, configure all paths in the HBA boot BIOS as primary boot device, secondary boot device, and so on. _____

Configuring SAN boot for Veritas Storage Foundation

Configuring SAN boot on the host

To configure SAN boot on the host using a storage system, complete the following steps.

Step	Action
1	On the storage system, select a supported cfmode.
2	Create a LUN on the storage system which acts as the SAN boot LUN.
3	Map the LUN to the host.
4	Ensure that only one non-proxy path is available to the LUN.
5	Ensure that only one SAN boot LUN is available to the host and that LUN is mapped as LUN ID 0.
6	Enable the boot BIOS of the HBA port to which the SAN boot LUN is mapped. For information about how to enable the boot BIOS, see the respective HBA vendor-specific documentation. Note_____ Enable the spinup delay option for the HBA port. This permits the HBA BIOS to scan for three minutes until it finds a boot device. _____
7	Note_____ Before installing the OS see the section on rootability in the <i>Veritas Volume Manager Administrator's Guide</i> for Linux that is shipped along with the software for partitioning information. While installing SUSE Linux Enterprise Server 10 Operating System, ensure that GRUB is installed in Master Boot Record (MBR) by going to Expert tab in software package selection screen during installation. _____
	After performing the appropriate changes to the HBA BIOS and ensuring the SAN boot LUN is visible, install the operating system on the SAN boot LUN.
8	After the operating system installation, reboot the host. As a result, the host boots from the SAN boot LUN on the storage system through a non-proxy path.

Step	Action
9	Install FC Linux Host Utilities 4.2. For information about installing the FC Linux Host Utilities 4.2, see “ Installing the FC Linux Host Utilities software ” on page 16.
10	Install the supported version of the drivers, if you are using external HBA drivers. For more information about the installation and configuration of the drivers, see Chapter 2, “ Installing HBAs and Software on the Host ,” on page 15.
11	Verify the HBA settings. For more information, see “ Verifying the HBA settings ” on page 63.
12	Install Veritas Storage Foundation, any relevant patches or fixes, and respective ASL and APM.
13	The vxdump restore daemon should be configured with an interval of 60. To configure the daemon, run the following command: vxdmptune dmp_restore_interval 60 On reboot this value comes into effect and remains persistent.
14	Encapsulate the root disk for use in VxVM. Any data on the disk is preserved by encapsulation when the disk is placed under VxVM control. To encapsulate the disk run the following command: vxdiskadm For detailed steps, see the section on encapsulating the disk in <i>Veritas Volume Manager Administrator's Guide</i> for Linux that is shipped along with the software.
15	Reboot the host after encapsulation.
16	Verify the encapsulation by issuing the following command: vxprint It displays the rootvolume and swapvolume under the corresponding disk group.
17	Configure the paths to the HBA boot BIOS as primary, secondary, tertiary, and so on, on the boot device. For more information, see the respective HBA vendor-specific documentation. Note _____ You must have non-proxy paths as primary paths in the HBA BIOS. _____

About this chapter

If you are using Veritas Storage Foundation for multipathing, you must install and configure the Symantec® Array Support Library (ASL) and Array Policy Module (APM) for storage systems. This chapter describes how to install and configure the ASL and APM and how these components affect multipathing and failover.

Note

To determine which version of the Veritas Storage Foundation and ASL to use, check the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Overview of the Array Support Library and Array Policy Module](#)” on page 134
- ◆ “[Installing the ASL and APM](#)” on page 135
- ◆ “[Removing the ASL and APM](#)” on page 137
- ◆ “[Obtaining multipathing information](#)” on page 139
- ◆ “[Understanding ASL error message types](#)” on page 143

Overview of the Array Support Library and Array Policy Module

What the Array Support Library is

The Array Support Library (ASL) is a Data ONTAP qualified library that provides information about storage array attributes and multipathing configuration to the Device Discovery Layer (DDL) and Dynamic Multipathing (DMP) components of Veritas Volume Manager (VxVM).

The DDL is a component of VxVM that discovers available enclosure information and the attributes of 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 ASL provides enclosure based naming. This means the name of the disk is based on the logical name of its enclosure, or disk array. The ASL provides specific vendor and model information to DMP and VxVM instead of referring to them as JBOD or raw devices.

Note

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

What the Array Policy Module is

The Array Policy Module (APM) is a kernel module that defines I/O error handling, failover path selection, and other failover behavior for a specific array. The NetApp APM is customized to optimize I/O error handling and failover path selection for the NetApp environment. When the ASL detects the NetApp storage array, it sets the array type to A/P-C-NETAPP. This instructs Veritas software to use the NetApp APM to handle I/O and path failures for the NetApp storage array.

Installing the ASL and APM

ASL and APM installation overview

Installing the ASL and APM involves the following tasks:

- ◆ Verifying that your configuration meets system requirements.
See the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.
- ◆ Obtaining the ASL and APM
- ◆ Installing the latest version of the ASL and APM

Obtaining the ASL and APM software

The ASL and APM are not distributed with the Host Utilities. You can obtain the ASL and APM from the Symantec Web site. For a direct link to the ASL and APM on the Symantec Web site, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml. Download this software before you start the installation.

Installing ASL and APM software you downloaded

To install the latest version of the ASL and APM that you downloaded from Symantec Web site, complete the following steps.

Step	Action
1	Log in to the system or host.
2	If you already have NetApp storage configured as JBOD in your VxVM configuration, remove the JBOD support for NetApp by entering the following command: vxddladm rmjbod vid=NETAPP
3	Install the ASL and APM according to the directions provided by Symantec.

Step	Action
4	If your host is connected to NetApp storage, you can verify the installation by locating the NetApp Enclosure Type in the output of the vxddmpadm listenclosure all command. The output shows the model name of the storage device if you are using enclosure-based naming with VxVM. vxddmpadm listenclosure all Example: The vxddmpadm listenclosure all command shows the Enclosure Type as FAS3020 in this example. <pre>[root@lnx132 ~]# vxddmpadm listenclosure all ENCLR_NAME ENCLR_TYPE ENCLR_SNO STATUS ARRAY_TYPE LUN_COUNT ===== disk Disk DISKS CONNECTED Disk 2 fas30200 FAS3020 1072458 CONNECTED A/A-NETAPP 102</pre>

Removing the ASL and APM

Before you begin

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

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

Removing the ASL and APM

To remove the ASL and APM, complete the following steps.

Step	Action
1	Log in to the host as the root user.
2	Stop all I/O to LUNs configured on storage controllers. Note_____ In a Storage Foundation RAC cluster, you must also stop clustering on a node before you remove the ASL and APM. _____

Step	Action
3	◆ Enter the following command to remove the ASL package: <code>rpm -ev <i>asl_rpm_name</i></code> Examples: The following command removes a previous version of the ASL on a Red Hat Enterprise Linux host. <code>rpm -ev VRTSNTAPas1-5.0-3.0_RHEL5</code> The following command removes a previous version of the ASL on a SUSE Linux Enterprise Server host. <code>rpm -ev VRTSNTAPas1-5.0-3.0_SLES10</code> ◆ Enter the following command to remove the APM package: <code>rpm -ev <i>apm_rpm_name</i></code> Example: The following command removes a previous version of the APM on a Red Hat Enterprise Linux host. <code>rpm -ev VRTSNTAPapm-5.0-1.0_RHEL5</code> The following command removes a previous version of the APM on a SUSE Linux Enterprise host. <code>rpm -ev VRTSNTAPapm-5.0-1.0_SLES10</code>

Obtaining multipathing information

What an ASL array type is

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

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

For additional information about system management, see the *Veritas Volume Manager Administrator's Guide*.

Using VxVM to display available paths

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

Step	Action																																																							
1	Enter the following command to view all devices: <pre>vxdisk list</pre> Result: The VxVM management interface displays the vxdisk device, type, disk, group, and status. It also shows which disks are managed by VxVM. Example: The following is a sample output of the vxdisk list command. <pre>[root@lnx132 ~]# vxdisk list</pre> <table><tr><th>DEVICE</th><th>TYPE</th><th>DISK</th><th>GROUP</th><th>STATUS</th></tr><tr><td>disk_0</td><td>auto:none</td><td>-</td><td>-</td><td>online invalid</td></tr><tr><td>disk_1</td><td>auto</td><td>-</td><td>-</td><td>error</td></tr><tr><td>fas30200_0</td><td>auto</td><td>-</td><td>-</td><td>error</td></tr><tr><td>fas30200_1</td><td>auto:cdsdisk</td><td>fas30200_1</td><td>sharedg</td><td>online shared</td></tr><tr><td>fas30200_2</td><td>auto:cdsdisk</td><td>fas30200_79</td><td>sharedg1_2</td><td>online shared</td></tr><tr><td>fas30200_3</td><td>auto:cdsdisk</td><td>fas30200_80</td><td>sharedg1_2</td><td>online shared</td></tr><tr><td>fas30200_4</td><td>auto:cdsdisk</td><td>fas30200_2</td><td>sharedg</td><td>online shared</td></tr><tr><td>fas30200_5</td><td>auto:cdsdisk</td><td>fas30200_3</td><td>sharedg</td><td>online shared</td></tr><tr><td>fas30200_6</td><td>auto:cdsdisk</td><td>fas30200_101</td><td>sharedg3_4</td><td>online shared</td></tr><tr><td>fas30200_7</td><td>auto:cdsdisk</td><td>fas30200_102</td><td>sharedg3_4</td><td>online shared</td></tr></table> Note For Veritas Storage Foundation 5.0 MP1 and MP2, the ASL displays the enclosure-based naming disk objects in uppercase. For Veritas Storage Foundation 5.0 MP3, the default behavior of the ASL is to display the enclosure-based naming disk objects in lower case. To change the enclosure names to uppercase, run the following command: <pre>vxddladm set namingscheme=ebn lowercase=no</pre>	DEVICE	TYPE	DISK	GROUP	STATUS	disk_0	auto:none	-	-	online invalid	disk_1	auto	-	-	error	fas30200_0	auto	-	-	error	fas30200_1	auto:cdsdisk	fas30200_1	sharedg	online shared	fas30200_2	auto:cdsdisk	fas30200_79	sharedg1_2	online shared	fas30200_3	auto:cdsdisk	fas30200_80	sharedg1_2	online shared	fas30200_4	auto:cdsdisk	fas30200_2	sharedg	online shared	fas30200_5	auto:cdsdisk	fas30200_3	sharedg	online shared	fas30200_6	auto:cdsdisk	fas30200_101	sharedg3_4	online shared	fas30200_7	auto:cdsdisk	fas30200_102	sharedg3_4	online shared
DEVICE	TYPE	DISK	GROUP	STATUS																																																				
disk_0	auto:none	-	-	online invalid																																																				
disk_1	auto	-	-	error																																																				
fas30200_0	auto	-	-	error																																																				
fas30200_1	auto:cdsdisk	fas30200_1	sharedg	online shared																																																				
fas30200_2	auto:cdsdisk	fas30200_79	sharedg1_2	online shared																																																				
fas30200_3	auto:cdsdisk	fas30200_80	sharedg1_2	online shared																																																				
fas30200_4	auto:cdsdisk	fas30200_2	sharedg	online shared																																																				
fas30200_5	auto:cdsdisk	fas30200_3	sharedg	online shared																																																				
fas30200_6	auto:cdsdisk	fas30200_101	sharedg3_4	online shared																																																				
fas30200_7	auto:cdsdisk	fas30200_102	sharedg3_4	online shared																																																				
2	On the host console, enter the following command to display path information for the device you want: <pre>vxdmpadm getsubpaths dmpnodename=device</pre> <i>device</i> is the name listed under the output of the vxdisk list command. Example: The following is sample output. <pre>[root@lnx162 ~]# vxddmpadm getsubpaths dmpnodename=fas30700_174</pre> <table><tr><th>NAME</th><th>STATE[A]</th><th>PATH-TYPE[M]</th><th>CTLR-NAME</th><th>ENCLR-TYPE</th><th>ENCLR-NAME</th><th>ATTRS</th></tr><tr><td colspan="7">=====</td></tr><tr><td>sdct</td><td>ENABLED(A)</td><td>PRIMARY</td><td>c0</td><td>FAS3070</td><td>fas30700</td><td>-</td></tr><tr><td>sdgp</td><td>ENABLED</td><td>SECONDARY</td><td>c0</td><td>FAS3070</td><td>fas30700</td><td>-</td></tr><tr><td>sdnk</td><td>ENABLED</td><td>SECONDARY</td><td>c1</td><td>FAS3070</td><td>fas30700</td><td>-</td></tr><tr><td>sdof</td><td>ENABLED(A)</td><td>PRIMARY</td><td>c1</td><td>FAS3070</td><td>fas30700</td><td>-</td></tr></table>	NAME	STATE[A]	PATH-TYPE[M]	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS	=====							sdct	ENABLED(A)	PRIMARY	c0	FAS3070	fas30700	-	sdgp	ENABLED	SECONDARY	c0	FAS3070	fas30700	-	sdnk	ENABLED	SECONDARY	c1	FAS3070	fas30700	-	sdof	ENABLED(A)	PRIMARY	c1	FAS3070	fas30700	-													
NAME	STATE[A]	PATH-TYPE[M]	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS																																																		
=====																																																								
sdct	ENABLED(A)	PRIMARY	c0	FAS3070	fas30700	-																																																		
sdgp	ENABLED	SECONDARY	c0	FAS3070	fas30700	-																																																		
sdnk	ENABLED	SECONDARY	c1	FAS3070	fas30700	-																																																		
sdof	ENABLED(A)	PRIMARY	c1	FAS3070	fas30700	-																																																		

Step	Action
3	To obtain path information for a host HBA, enter the following command: <code>vxdmpadm getsubpaths ctrl=controller_name</code> <i>controller_name</i> is the controller displayed under “CTRL-NAME” in the output of the <code>vxdmpadm getsubpaths dmpnodename</code> command. Result: The output displays information about the paths to the storage system (whether the path is a primary or secondary path). The output also lists the storage system that the device is mapped to. Example: The following is sample output. <pre>[root@lnx162 ~]# vxdmpadm getsubpaths ctrl=c0 NAME STATE[A] PATH-TYPE[M] DMPNODENAME ENCLR-TYPE ENCLR-NAME ATTRS ===== sdbx DISABLED PRIMARY fas30700_100 FAS3070 fas30700 - sdft ENABLED(A) SECONDARY fas30700_100 FAS3070 fas30700 - sdbh DISABLED PRIMARY fas30700_101 FAS3070 fas30700 - sdfd ENABLED(A) SECONDARY fas30700_101 FAS3070 fas30700 -</pre>

Using sanlun to obtain multipathing information

To obtain multipathing information by using the **sanlun** command, complete the following step.

Step	Action																																																																																																								
1	On the host command line, enter the following command: <pre>sanlun lun show -p</pre> Result: sanlun displays path information for each LUN. Example: The following is sample output of the sanlun lun show -p command: <pre>f3040-200-90:/vol/vol21/lun21 (LUN 23) Lun state: GOOD Lun Size: 2g (2147483648) Filer_CF_State: Cluster Enabled Protocol: FCP Filer Partner: f3040-200-89 DMP Nodename: fas30400_112 Multipath-provider: VERITAS</pre> <table><tr><td>Veritas</td><td>sanlun</td><td>Filer</td><td></td><td></td><td>Primary</td><td>Partner</td><td>Veritas</td></tr><tr><td>path</td><td>path</td><td>Path</td><td>/dev/</td><td>Host</td><td>Filer</td><td>Filer</td><td>IO</td></tr><tr><td>state</td><td>state</td><td>type</td><td>node</td><td>HBA</td><td>port</td><td>port</td><td>Policy</td></tr><tr><td colspan="8">-----</td></tr><tr><td>-</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>ENABLED</td><td>GOOD</td><td>SECONDARY</td><td>sdk</td><td>host3</td><td>--</td><td>0b</td><td>Round-Robin</td></tr><tr><td>ENABLED(A)</td><td>GOOD</td><td>PRIMARY</td><td>sdj</td><td>host3</td><td>0b</td><td>--</td><td>Round-Robin</td></tr><tr><td>ENABLED(A)</td><td>GOOD</td><td>PRIMARY</td><td>sdh</td><td>host3</td><td>0d</td><td>--</td><td>Round-Robin</td></tr><tr><td>ENABLED</td><td>GOOD</td><td>SECONDARY</td><td>sdg</td><td>host3</td><td>--</td><td>0d</td><td>Round-Robin</td></tr><tr><td>ENABLED</td><td>GOOD</td><td>SECONDAR</td><td>sdz</td><td>host2</td><td>--</td><td>0c</td><td>Round-Robin</td></tr><tr><td>ENABLED(A)</td><td>GOOD</td><td>PRIMARY</td><td>sdep</td><td>host2</td><td>0a</td><td>--</td><td>Round-Robin</td></tr><tr><td>ENABLED</td><td>GOOD</td><td>SECONDARY</td><td>sddb</td><td>host2</td><td>--</td><td>0a</td><td>Round-Robin</td></tr><tr><td>ENABLED(A)</td><td>GOOD</td><td>PRIMARY</td><td>sdbn</td><td>host2</td><td>0c</td><td>--</td><td>Round-Robin</td></tr></table>	Veritas	sanlun	Filer			Primary	Partner	Veritas	path	path	Path	/dev/	Host	Filer	Filer	IO	state	state	type	node	HBA	port	port	Policy	-----								-								ENABLED	GOOD	SECONDARY	sdk	host3	--	0b	Round-Robin	ENABLED(A)	GOOD	PRIMARY	sdj	host3	0b	--	Round-Robin	ENABLED(A)	GOOD	PRIMARY	sdh	host3	0d	--	Round-Robin	ENABLED	GOOD	SECONDARY	sdg	host3	--	0d	Round-Robin	ENABLED	GOOD	SECONDAR	sdz	host2	--	0c	Round-Robin	ENABLED(A)	GOOD	PRIMARY	sdep	host2	0a	--	Round-Robin	ENABLED	GOOD	SECONDARY	sddb	host2	--	0a	Round-Robin	ENABLED(A)	GOOD	PRIMARY	sdbn	host2	0c	--	Round-Robin
Veritas	sanlun	Filer			Primary	Partner	Veritas																																																																																																		
path	path	Path	/dev/	Host	Filer	Filer	IO																																																																																																		
state	state	type	node	HBA	port	port	Policy																																																																																																		

-																																																																																																									
ENABLED	GOOD	SECONDARY	sdk	host3	--	0b	Round-Robin																																																																																																		
ENABLED(A)	GOOD	PRIMARY	sdj	host3	0b	--	Round-Robin																																																																																																		
ENABLED(A)	GOOD	PRIMARY	sdh	host3	0d	--	Round-Robin																																																																																																		
ENABLED	GOOD	SECONDARY	sdg	host3	--	0d	Round-Robin																																																																																																		
ENABLED	GOOD	SECONDAR	sdz	host2	--	0c	Round-Robin																																																																																																		
ENABLED(A)	GOOD	PRIMARY	sdep	host2	0a	--	Round-Robin																																																																																																		
ENABLED	GOOD	SECONDARY	sddb	host2	--	0a	Round-Robin																																																																																																		
ENABLED(A)	GOOD	PRIMARY	sdbn	host2	0c	--	Round-Robin																																																																																																		

Understanding ASL error message types

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

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

Troubleshooting

Note

For a list of known problems and limitations, be sure to check the *Release Notes* for the Host Utilities.

Blacklisting the devices

To blacklist the devices, complete the following steps.

Step	Action
1	Identify the entries in the sysfs that are not displayed by the sanlun lun show command.
2	Blacklist those sd entries.

Unloading of the QLogic module fails

Problem: On the Linux host, unloading the QLogic module sometimes fails.

When Linux FC initiators are included in multiple initiator FC switch zones, the HBA module usage count on the specified initiator host always shows a higher value than expected despite no apparent usage of the HBA module.

For example, for a SUSE Linux Enterprise Server 9 SP3 host with its QLogic initiators connected to a multiple initiator zone on the FC switch, the QLogic module usage count is sometimes shown as follows:

```
qla2400 207360 0
qla2xxx 3654906 5 qla2400
```

Here, the qla2xxx usage count is shown as 5 (instead of 1), despite the QLogic module not being used. This can cause unloading of the QLogic module to fail.

Workaround: While connecting FC Linux hosts to a switch fabric, always ensure that each FC host initiator is connected to a separate individual initiator zone alone instead of a multiple initiator zone.

Linux FC hosts with QLogic initiators fail to detect LUNs

Problem: Linux FC hosts with QLogic initiators fail to detect LUNs mapped to it when configured to auto negotiate port topology and speed. In hosts using QLogic initiators, the host initiator fails to establish a consistent topology and speed, despite setting both to auto negotiate.

This can cause the host to miss a few LUNs that are mapped to it, during the HBA scan. This behavior persists even after rebooting the host.

Workaround: When using QLogic initiators on a Linux FC host, ensure the following settings are done for port topology and speed using the SANsurfer scli package as follows:

1. Set the value of QLogic HBA port topology to Point to Point Only while connecting to a NetApp storage controller operating in SSI, Standby or Partner CF modes, or to Loop Only if the storage system is operating in Dual Fabric or Mixed CF modes.

To do this, perform the following steps:

- a. Run **/opt/QLogic_Corporation/SANsurferCLI/scli**.
- b. Select Configure HBA Settings from the Main Menu.
- c. For each of the WWPNs listed here, select ConnectionOptions and set it to Point to Point Only or Loop Only as required.

2. Set the value of QLogic port speed ideally to the highest speed possible, depending on its ceiling and that of the switch/target port it is connected to.

To do this, perform the following steps:

- a. Run the **/opt/QLogic_Corporation/SANsurferCLI/scli**.
- b. Select the Configure HBA Settings from the Main Menu.

3. For each of the WWPNs listed here, select the Data Rate option and set it to the specified speed.

LUNs are not visible on the host

FC LUNs appear as local disks to the host. If the storage system LUNs are not available as disks on the hosts, try rescanning the HBAs by reloading the drivers.

If the LUNs still do not appear, verify the following configuration settings.

Configuration setting	What to do
Driver version	Verify that the external HBA driver version you are running is supported. You can check the most up-to-date system requirements in the NetApp FC Host Compatibility Matrix at the following URL: http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/
Cabling	Verify that the cables between the host and the storage system are connected.
Zoning	Verify that the zoning is correct. For information about zoning see “ Switched fabric configurations ” on page 48.
HBA Setting	Verify the recommended HBA settings. For information about verifying the HBA settings, see “ Verifying the HBA settings ” on page 63.
FC service status	Verify that the FC service is licensed and started on the storage system according to the procedure described in the <i>Data ONTAP Block Access Management Guide</i> .
Initiator login	Verify that the initiator is logged in to the storage system by entering the fcp show initiator 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 WWPN.
HBA World Wide Port Names	Verify that you are using the correct initiator WWPNs in the igroup configuration. On the storage system, use the igroup show command to display the WWPNs of the initiators in the storage system’s igroups. On the host, use the HBA tools to display the initiator WWPN. The initiator WWPNs configured in the igroup and on the host must match.

Configuration setting	What to do
LUN mappings	Verify that the LUNs are mapped to an igroup. On the storage system, use one of the following commands: ◆ lun show -m displays all LUNs and the igroups to which they are mapped. ◆ lun show -g igroup-name displays the LUNs mapped to a specific igroup.
System requirements	Verify that the components of your configuration are supported. Verify that you have the correct host Operating System (OS) service pack level, initiator version, Data ONTAP version, and other system requirements. You can check the most up-to-date system requirements in the NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

LUNs not recognized by dm-multipath

Problem: LUNs are mapped to the Linux host and are correctly reported by the **sanlun lun show all** command. However, the LUNs are not reported by the **multipath -ll** command. Multipath devices are not created for the LUNs.

Workaround: Ensure the LUNs are not excluded (blacklisted) from the dm-multipath configuration. Enter the following command on the Linux host to display the devices that are blacklisted:

```
multipath -v -d | grep blacklist
```

If any FC devices are blacklisted, check `devnode_blacklist` or the `blacklist` section of the `/etc/multipath.conf` file. Ensure that the entries are correctly specified.

For information about blacklisting devices on a Red Hat Enterprise Linux host, see “[Configuring dm-multipath](#)” on page 55.

For more information about blacklisting devices on a SUSE Linux Enterprise Server host, see “[Configuring dm-multipath](#)” on page 86.

If the FC devices are not blacklisted, but are still not recognized by the multipath command, regenerate the multipath maps by entering the following command:

```
multipath -v3
```

For more information, see bug 228744 on Bugs Online at <http://now.netapp.com/cgi-bin/bol>.

FC performance is poor

Problem: If you experience poor performance on some or all FC LUNs that have partitions, the alignment between the LUN block boundaries and the underlying WAFL boundaries might not match. This problem occurs only with partitions and only with certain combinations of disk geometry parameters.

Workaround: Manually enter the disk geometry parameters when creating partitions with the `fdisk` command. For specific steps, refer to bug number 156121 on Bugs Online at <http://now.netapp.com/NOW/cgi-bin/bol>.

Error message from `sanlun fcp show adapter`

Problem: The `sanlun fcp show` adapter command returns the following error message:

Unable to locate /usr/lib/libHBAAPI.so library

Ensure the package installing the library is installed & loaded

Workaround: The libHBAAPI.so file is included with the QLogic SANsurfer Pro package and the Emulex applications package. Install the appropriate package for your HBA model. See “[Downloading and installing the external HBA driver and applications](#)” on page 22.

LUNs unavailable after storage system failover or giveback

Problem: The host requires specific driver, HBA, and Fibre Channel switch zoning settings to maintain LUN access across storage system takeover, giveback, and reboot.

Workaround: Verify that you have set the required driver and HBA settings. Verify that you have correctly zoned your Fibre Channel switches.

For multipath LUNs, you can restore access to the LUNs by restarting the multipathing daemon and updating the multipathing maps by running the following commands:

```
/etc/init.d/multipathd start
```

```
multipath -v3
```

**sanlun commands
fail on SLES 9
x86_64 hosts**

Problem: `sanlun` commands fail on SLES 9 x86_64 hosts.

Workaround: In SUSE Linux Enterprise Server 9 x86_64 hosts, there is no 32-bit compatibility library available in the installation media for `libdevmapper`. Because `sanlun` is a 32-bit application, it can link to only 32-bit libraries. When `sanlun` is run on such a machine, it exits with the error message that it cannot find the `libdevmapper.so` library.

Download the 32-bit compatibility rpm for `device-mapper-multipath` from the Novell Web site at <http://support.novell.com/techcenter/psdb/8ae2a27b3b63ac95fd776298026596eb.html>. Install the rpm which provides the 32-bit `libdevmapper.so` compatibility library. When `sanlun` is run, it links to the 32-bit library.

Troubleshooting—Veritas Storage Foundation

Note

Be sure to check the *Release Notes* for the Host Utilities for a list of known problems and limitations.

Error when stopping the fencing driver

When you get an error stopping the fencing driver, complete the following steps.

Step	Action
1	Stop the Veritas cluster service using <code>hastop</code> on all the nodes.
2	Try stopping the fencing driver again.
3	If it does not work, remove the name of the coordinator diskgroup from <code>/etc/vxfendg</code> .
4	In the <code>/etc/vxfenmode</code> file, make the following change: <code>vxfen_mode=disabled</code>
5	Stop the fencing driver by using the following command: <code>/etc/init.d/vxfen stop</code>

Note

It is recommended to halt all the nodes except the last one in the cluster.

Starting the vxconfigd

To start the `vxconfigd` debugging at debug level 5, run the following command:

```
vxconfigd -k -x5 -x log
```

The corresponding log file where `vxclnt` daemon stores the debug messages is `/etc/vx/vxconfigd.log`.

To start `vxconfigd` again without debugging as earlier, run the following command:

```
vxconfigd -k -x cleartempdir -x syslog
```

Setting the dmp_log_level

To set `dmp_log_level` to 4 (highest debugging level), run the following command:

```
vxddmpadm settune dmp_log_level=4
```

Secondary paths identified incorrectly

Problem: Secondary paths identified incorrectly as primary paths while configuring DMP.

Workaround: Verify whether Array Support Library (ASL) for NetApp storage is installed correctly. That is, the ASL is successfully installed without any error messages.

Enclosure based naming not reflected on NetApp storage

Problem: VxVM's enclosure based naming feature is enabled, but the DMP devices do not contain a NetApp storage model string.

Workaround: Verify whether the Array Support Library (ASL) for NetApp storage is installed correctly. That is, the ASL is successfully installed without any error messages.

SAN boot installation to LUN configured on storage controller fails

Verify whether the host can only see one path to the LUN configured on storage controller during installation.

Encapsulation of root disk fails after SAN boot installation

Make sure to follow the partitioning scheme given in the *Veritas™ Volume Manager Administrator's Guide* for Linux.

About this information

The following copyright statements and licenses apply to open source software components that are distributed with various versions of the NetApp Host Utilities software products. Your product does not necessarily use all the open source software components referred to below.

Where required, source code is published at the following location:

<ftp://ftp.netapp.com/frm-ntap/opensource/>

Copyrights and licenses

The following components are subject to the Perl Artistic License:

- ◆ FTP.pm portions copyright © 1995-2004 Graham Barr <gbarr@pobox.com>. All rights reserved.
- ◆ Telnet.pm portions copyright © 1997, 2000, 2002 Jay Rogers. All rights reserved.

Perl Artistic License:

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

Definitions:

“Package” refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

“Standard Version” refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder.

“Copyright Holder” is whoever is named in the copyright or copyrights for the package.

“You” is you, if you’re thinking about copying or distributing this Package.

“Reasonable copying fee” is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

“Freely Available” means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.
2. You may apply bug fixes, portability fixes, and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:
 - Place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as ftp.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.
 - Use the modified Package only within your corporation or organization.
 - Rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.
 - Make other distribution arrangements with the Copyright Holder.
4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:
 - Distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.
 - Accompany the distribution with the machine-readable source of the Package with your modifications.
 - Accompany any non-standard executables with their corresponding Standard Version executables, giving the non-standard executables non-standard names, and clearly documenting the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.
 - Make other distribution arrangements with the Copyright Holder.
5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package.

You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own.

6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whomever generated them, and may be sold commercially, and may be aggregated with this Package.
7. C or perl subroutines supplied by you and linked into this Package shall not be considered part of this Package.
8. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.
9. THIS PACKAGE IS PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The HBA API component is subject to the SNIA Public License 1.1.

Portions copyright © 2002 International Business Machines Corporation. All rights reserved.

Portions authored by Benjamin F. Kuo, Troika Networks, Inc.

Portions copyright © 2001 Microsoft Corporation. All rights reserved.

SNIA Public License 1.1:

1. Definitions.

1.1 “Commercial Use” means distribution or otherwise making the Covered Code available to a third party.

1.2 “Contributor” means each entity that creates or contributes to the creation of Modifications.

1.3 “Contributor Version” means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.4 “Covered Code” means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.5 “Electronic Distribution Mechanism” means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.6 “Executable” means Covered Code in any form other than Source Code.

1.7 “Initial Developer” means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.8 “Larger Work” means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.9 “License” means this document.

1.10 “Licensable” means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.11 “Modifications” means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

- ❖ Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.
- ❖ Any new file that contains any part of the Original Code or previous Modifications.

1.12 “Original Code” means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.13 “Patent Claims” means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.14 “Source Code” means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor’s choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.15 “You” (or “Your”) means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, “You” includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, “control” means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty

percent (50%) of the outstanding shares or beneficial ownership of such entity

2. Source Code License.

2.1 The Initial Developer Grant. The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third-party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense, and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patents Claims infringed by the making, using, or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

2.2 Contributor Grant. Subject to third-party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense, and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2)

separate from the Contributor Version; 3) for infringements caused by: i) third-party modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1 Application of License. The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

3.2 Availability of Source Code. Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

3.3 Description of Modifications. You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

3.4 Intellectual Property Matters.

(a) Third Party Claims. If Contributor has actual knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail

that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter.

(b) Contributor APIs. If Contributor's Modifications include an application programming interface and Contributor has actual knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations. Contributor represents that, except as disclosed pursuant to Section 3.4(a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

3.5 Required Notices. You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be most likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity, or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability (excluding any liability arising from intellectual property claims relating to the Covered Code) incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity, or liability terms You offer.

3.6 Distribution of Executable Versions. You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligation of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a

different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability (excluding any liability arising from intellectual property claims relating to the Covered Code) incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.7 Larger Works. You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation. If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the LEGAL file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.
5. Application of this License. This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.
6. Versions of the License.

6.1 New Versions. The Storage Networking Industry Association (the “SNIA”) may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2 Effect of New Versions. Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by the SNIA. No one other than the SNIA has the right to modify the terms applicable to Covered Code created under this License.

6.3 Derivative Works. If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases “Storage Networking Industry Association,” “SNIA,” or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the SNIA Public License. (Filling in the name of the Initial Developer, Original Code,

or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

7. DISCLAIMER OF WARRANTY. COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN “AS IS” BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR, OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

8. TERMINATION.

8.1 This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within a reasonable time after becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2 If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as “Participant”) alleging that:

(a) such Participant’s Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

8.3 If You assert a patent infringement claim against Participant alleging that such Participant’s Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior

to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4 In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

9. **LIMITATION OF LIABILITY.** UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.
10. **U.S. GOVERNMENT END USERS.** The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.
11. **MISCELLANEOUS.** This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation

which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

12. **RESPONSIBILITY FOR CLAIMS.** As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.
13. **MULTIPLE-LICENSED CODE.** Initial Developer may designate portions of the Covered Code as “Multiple-Licensed”. “Multiple-Licensed” means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of this License or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.
14. **ACCEPTANCE.** This License is accepted by You if You retain, use, or distribute the Covered Code for any purpose.

EXHIBIT A—The SNIA Public License.

The contents of this file are subject to the SNIA Public License Version 1.1 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

This file itself is the license.

Software distributed under the License is distributed on an “AS IS” basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is _____.

The Initial Developer of the Original Code is _____.

Contributor(s): _____.

The Net_SSLeay.pm component is subject to the Net SSLeay License.

NetSSLeay License:

http://search.cpan.org/src/SAMPO/Net_SSLeay.pm-1.25/README License and Copying: Copyright © 1996-2002 Sampo Kellomaki. All rights reserved.

Distribution and use of this module is under the same terms as the OpenSSL package itself (i.e. free, but mandatory attribution; NO WARRANTY). Please consult LICENSE file in the root of the OpenSSL distribution. While the source distribution of this perl module does not contain Eric’s or OpenSSL’s code, if you use this module you will use OpenSSL library. Please give Eric and OpenSSL team credit (as required by their licenses). And remember, you, and

nobody else but you, are responsible for auditing this module and OpenSSL library for security problems, backdoors, and general suitability for your application.

SSLey License copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscape's SSL. This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)". The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related.
4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgment: "This product includes software written by Tim Hudson (tjh@cryptsoft.com)".

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS 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. The licence and distribution terms for any publicly available version or derivative of this code cannot be changed; i.e., this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.

The Netkit component is subject to the GNU GPL 2.0 license.

Copyright © 1989, 1991 Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

GNU General Public License 2.0:

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation’s software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution, and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION, AND MODIFICATION

This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
 - c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute, or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients'

exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement, or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this

License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found:

One line to give the program’s name and a brief idea of what it does. Copyright (c)

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w`. This is free software, and you are welcome to redistribute it under certain conditions; type `show c` for details.

The hypothetical commands `show w` and `show c` should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w` and `show c`; they could even be mouse-clicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program ‘Gnomovision’ (which makes passes at compilers) written by James Hacker.

signature of Ty Coon, 1 April 1989

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

The Open Hardware Certification component is subject to the Robin's Nest Software license.

Copyright © 1990-2008 Robin's Nest Software, Inc. All rights reserved.

Robin's Nest Software License:

This Software Provided By Robin's Nest Software Inc.

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in the supporting documentation, and that the name of the author not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

THE AUTHOR DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL HE BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE, OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

The Windows Installer XML (WiX) Toolset component is subject to the Common Public License.

Copyright © Microsoft Corporation. All rights reserved.

Common Public License:

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS COMMON PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION, OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPIENT'S ACCEPTANCE OF THIS AGREEMENT.

1. DEFINITIONS

"Contribution" means:

- a) in the case of the initial Contributor, the initial code and documentation distributed under this Agreement, and

b) in the case of each subsequent Contributor:

- i) changes to the Program, and
- ii) additions to the Program;

where such changes and/or additions to the Program originate from and are distributed by that particular Contributor. A Contribution ‘originates’ from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor’s behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

“Contributor” means any person or entity that distributes the Program.

“Licensed Patents” mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

“Program” means the Contributions distributed in accordance with this Agreement.

“Recipient” means anyone who receives the Program under this Agreement, including all Contributors.

2. GRANT OF RIGHTS

a) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

b) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import, and otherwise transfer the Contribution of such Contributor, if any, in source code and object code form. This patent license shall apply to the combination of the Contribution and the Program if, at the time the Contribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations which include the Contribution. No hardware per se is licensed hereunder.

c) Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the

rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third-party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibility to acquire that license before distributing the Program.

d) Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

3. REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

- a) it complies with the terms and conditions of this Agreement; and
- b) its license agreement:
 - i) effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;
 - ii) effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, special, incidental, and consequential damages, such as lost profits;
 - iii) states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and
 - iv) states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

- a) it must be made available under this Agreement; and
- b) a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reasonably allows subsequent Recipients to identify the originator of the Contribution.

4. COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, business partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor

who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore, if a Contributor includes the Program in a commercial product offering, such Contributor (“Commercial Contributor”) hereby agrees to defend and indemnify every other Contributor (“Indemnified Contributor”) against any losses, damages, and costs (collectively “Losses”) arising from claims, lawsuits, and other legal actions brought by a third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offering. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must: a) promptly notify the Commercial Contributor in writing of such claim, and b) allow the Commercial Contributor to control, and cooperate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense.

For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Commercial Contributor’s responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

5. NO WARRANTY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of using and distributing the Program and assumes all risks associated with its exercise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

6. DISCLAIMER OF LIABILITY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL,

SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), 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 OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. GENERAL

If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

If Recipient institutes patent litigation against a Contributor with respect to a patent applicable to software (including a cross-claim or counterclaim in a lawsuit), then any patent licenses granted by that Contributor to such Recipient under this Agreement shall terminate as of the date such litigation is filed. In addition, if Recipient institutes patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipient agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Program shall continue and survive.

Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsistency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify this Agreement. IBM is the initial Agreement Steward. IBM may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published,

Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or otherwise. All rights in the Program not expressly granted under this Agreement are reserved.

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting litigation.

