FCP Linux® Host Utilities 4.1 for Native OS and Veritas®
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 215-03802_A0
April 2008

Copyright and trademark information

Copyright information

Copyright © 1994–2008 NetApp, Inc. All rights reserved. Printed in the U.S.A.

No part of this file 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 material of NetApp, Inc. 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 set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, Data ONTAP, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, NOW NetApp on the Web, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, StoreVault, SyncMirror, Topio, VFM, and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. Cryptainer, Cryptoshred, Datafort, and Decru are registered trademarks, and Lifetime Key Management and OpenKey are trademarks, of Decru, a NetApp, Inc. company, in the U.S.A. and/or other countries. SANScreen is a registered trademark of Onaro, Inc., a NetApp, Inc. company, 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; LockVault; NOW; ONTAPI; RAID-DP; ReplicatorX; RoboCache; RoboFiler; SecureAdmin; Serving Data by Design; SharedStorage; Simplicore; Simulate ONTAP; Smart SAN; SnapCache; SnapDirector; SnapFilter;

SnapMigrator; SnapSuite; SohoFiler; SpinMirror; SpinRestore; SpinShot; SpinStor; vFiler; VFM Virtual File Manager; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, Inc. in the U.S.A.

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

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

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

Table of Contents

Preface ix

Chapter 1 **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 FCP Linux Host Utilities software 16

Removing the previous 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 26

Recording the WWPNs of the HBAs 29

Chapter 3 **Configuring the Storage System for FCP** 31

Chapter 4 **Migrating a storage system to single_image mode** 33

Understanding cfmode 34

Migrating to single_image mode 36

Chapter 5 **Configuring FCP cabling and switches** 41

Switched fabric configurations 42

Direct-connected configurations 44

Chapter 6 **Accessing LUNs from a Red Hat Enterprise Linux host** 45

Loading the HBA driver on the host 47

Loading the bundled HBA driver on the host 48

Configuring the dm-multipath 49

Starting the multipath service. 53

Accessing LUNs with dm-multipath support. 54
- Verifying the HBA setting and dm-multipath configuration. 55
- Identifying and viewing a list of LUNs 58
- Creating and mounting a file system on a multipath device 60
- Tuning the dm-multipath configuration 62

Stopping the dm-multipath 63

Accessing LUNs without dm-multipath support 64
- Identifying and viewing LUN information. 65
- Creating a partition 66
- Creating a file system 68
- Labeling and mounting the file system. 69

Unloading the HBA driver and shutting down 71

Unloading the bundled HBA driver and shutting down 73

Chapter 7 **Accessing LUNs from a SUSE Linux Enterprise Server host** 75

Loading the HBA driver on the host 76

Configuring the dm-multipath 77

Starting the multipath service. 81

Accessing LUNs with dm-multipath support. 82
- Verifying the HBA setting, dm-multipath configuration 83
- Identifying and viewing a list of LUNs 86
- Creating and mounting a file system on a LUN 88
- Tuning the dm-multipath configuration 90

Stopping dm-multipath 91

Accessing LUNs without dm-multipath support 92
- Identifying and viewing the LUN information. 93
- Creating a partition 95
- Creating a file system 96
- Labelling and mounting the file system 97

Unloading the HBA driver and shutting down 99

Chapter 8 **Configuring SAN Boot on Red Hat Enterprise Linux host** 101
Configuring SAN boot on native host 102
Configuring SAN boot for Veritas Storage Foundation 108

Chapter 9 **Veritas Array Support Library and Array Policy Module** 113
Overview of the Array Support Library and Array Policy Module 114
Installing the ASL and APM 115
Removing the ASL and APM 119
Obtaining multipathing information 120
Understanding ASL error message types 123

Appendix A **Troubleshooting** 125

Appendix B **Troubleshooting — Veritas Storage Foundation** 131

Preface

About this guide

This guide describes how to install the FCP Linux® Host Utilities 4.1 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 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
Italic type	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the `fctest` *adaptername* command, you enter the characters "fctest" 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.
Enter	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards.
type	Used to mean pressing one or more keys on the keyboard.
enter	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.

Overview 1

About this chapter

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

About Linux versions

FCP 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. 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

FCP 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 FCP hosts.

The term *external HBA driver* is used to specify the drivers that is customized for use with a NetApp storage system and which 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.

This guide has setup procedures for both external and bundled HBA drivers. When referring to bundled drivers they are specifically called out as *bundled drivers* in this guide. If not specifically mentioned, the information in this guide applies to the external drivers. That is, the term *HBA driver* is used to refer the external drivers.

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 NetApp LUN, 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 FCP Linux Host Utilities is

FCP Linux Host Utilities 4.1 enables a Linux host to connect to your storage systems. FCP 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

FCP Linux Host Utilities 4.1 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
- An updated mpath_prio_ontap and mpath_prio_ontap.static files are included in the Host Utilities
- The `install` and `uninstall` command scripts to add and remove the FCP Linux Host Utilities components
- The following diagnostic scripts for collecting troubleshooting information:
  - linux_info - collects configuration information about your Linux host
  - 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
  - mcdata_info - collects information about McData® switches installed in the network
  - libHBAAPI.so - support library for sanlun
  - 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

Multipathing

Linux device mapper multipathing (dm-multipath) is supported on NetApp 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, FCP 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 multipathing

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 FCP 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 FCP 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 FCP 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 33 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 FCP for SCSI to access data on storage systems that run supported versions of Data ONTAP software.

The storage system is an FCP target device. Each host HBA is an initiator that uses FCP 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 FCP connection from the HBA to the storage system through a single switch

    A host is cabled to a single FCP switch that is connected by cable to redundant FCP ports on a high-availability storage system (active/active configuration). A fabric-attached single-path host has one HBA.
  - A dual-host FCP 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 FCP connection from the host HBA to stand-alone or high-availability (active/active) storage system configurations.

Note

Use redundant configurations with two FCP switches for high availability in production environments. However, direct FCP 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 Protocol (FCP) 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 HBA drivers ◆ Downloading and installing the supported HBA driver ◆ Installing the FCP Linux Host Utilities software ◆ Setting required HBA parameters ◆ Installing dm-multipath packages ◆ Recording the worldwide port names (WWPNs) of the HBAs ◆ Installing Veritas Storage Foundation, if you want to use the Veritas multipathing solution

Step	Action
	Note For information on installing Veritas Storage Foundation, see the *Veritas Storage Foundation™ Installation Guide* for Linux that is shipped along with the software. ❖ Install Array Support Library (ASL) and Array Policy Module (APM) See Chapter 9, "Veritas Array Support Library and Array Policy Module," on page 113. 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 FCP service ◆ Creating LUNs, creating an initiator group (igroup) that contains the host's FCP 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 FCP," on page 31.
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 33.
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 FCP cabling and switches," on page 41.

Step	Action
5	Discover and access the LUNs from the host, including the following: ◆ Loading the HBA driver modules ◆ Configuring multipathing using dm-multipath See Chapter 6, "Accessing LUNs from a Red Hat Enterprise Linux host," on page 45, or Chapter 7, "Accessing LUNs from a SUSE Linux Enterprise Server host," on page 75.

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 *Release Notes* 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 *System Configuration Guide* at https://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/.
Configuring the storage system	◆ The *Best Practices for Reliability* at http://now.netapp.com/NOW/knowledge/docs/docs.cgi. ◆ The Data ONTAP *Software Setup Guide* 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 *Installation and Administration Guide* for that version of SnapDrive for UNIX.
Managing SAN storage on the storage system	◆ The *Command Reference* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ The Data ONTAP *Block Access Management Guide* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ The FilerView online Help ◆ The *Best Practices for Reliability: Ongoing Maintenance* at http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm.
Installing Veritas Storage Foundation	The *Veritas Storage Foundation™ Installation Guide* for Linux that is shipped along with the software.
Release information about the products in the Veritas Storage Foundation Linux product line	The *Veritas Storage Foundation™ Release Notes* for Linux that is shipped along with the software.

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

If you need information about…	See…
Enabling ALUA	Data ONTAP® Block Access Management Guide for FCP 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 2

Configuration overview

This chapter explains how to install the FCP Linux Host Utilities, HBAs, and drivers, and how to complete the initial configuration of the Linux host. You will 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.

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 FCP Linux Host Utilities software" on page 16
- "Removing the previous 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 26
- "Recording the WWPNs of the HBAs" on page 29

Requirement for root privileges

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

Installing the FCP Linux Host Utilities software

Verifying the correct version

Verify that you have the correct version of FCP 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 FCP Linux Host Utilities file (netapp_linux_host_utilities_4_1.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 Novell's website at http://support.novell.com/techcenter/psdb/8ae2a27b3b63ac95fd776298026596eb.html. This package needs to be installed, to ensure that sanlun work 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 FCP 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_1.tar.gz`

Step	Action
4	To extract the files, enter the following command: **`tar -xvf netapp_linux_host_utilities_4_1.tar`**
5	Change to the netapp_linux_host_utilities_4_1 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 HBA drivers

Removing existing QLogic HBA drivers

You need to remove any existing HBA driver modules. To remove the QLogic driver modules, perform 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: **`pkill qlremote`**
3	Change to the directory where the driver installer files are located.
4	Uninstall the existing 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 alone). ◆ For a SANsurfer package that is currently 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`**

Removing existing Emulex HBA drivers

You need to remove any existing HBA driver modules. To remove the Emulex driver modules, perform 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 existing 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 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 using the appropriate method for your version of Linux.

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 NetApp storage system. Be sure to verify the required driver version 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 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. Locate 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. Locate 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.

A reboot is required after installing the HBA driver package.

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.

Installing HBA management software

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

QLogic: The QLogic management CLI (that is, the scli) is installed by the driver installation program. The SANsurfer FC HBA Manager should be downloaded and installed from the QLogic Web site.

Emulex: The management software is included in the Emulex Linux application kit.

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, complete the following steps on the Linux host using the QLogic `scli` command.

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

For Emulex HBAs

No special settings are required for supported Emulex HBAs. Default settings works fine in the case of Emulex.

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: The SANsurfer CLI for Linux should be downloaded and installed from the QLogic Web site.

Emulex: The Emulex Linux application kit should be downloaded and installed from the Emulex Web site. 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 `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 NetApp 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.

Step	Action	
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	**IF...**	**THEN...**
	You are using Red Hat Enterprise Linux 4	**a.** For `sanlun fcp show adapter` command to work, load a QLogic ioctl module *qioctlmod* separately. To load the module, enter the following command: **`modprobe -v qioctlmod`** **b.** In the modprobe.conf file, append the following line: **`options qla2xxx qlport_down_retry=30 ql2xsuspendcount=6`** **c.** Recreate the initrd by running the following command: **Note** Take a backup of the current initrd before running the `mkinitrd` command. This is recommended as a precaution. **`mkinitrd -f /boot/initrd-"`uname -r`".img `uname -r` --with=qioctlmod`** **d.** 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.

Step	Action	
		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`
	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 works fine. **Note** There are no qla2300/qla2400 firmware modules in the Red Hat Enterprise Linux 5 series. These modules have been integrated into the base qla2xxx module itself.

Emulex HBA and bundled driver parameters: No special settings are required for supported bundled Emulex HBAs. The default settings work fine in the case of Emulex.

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 FCP 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 using the SANsurfer FC HBA CLI client, complete the following steps.

Step	Action
1	Enter the following command on the Linux host console: `scli -i`
2	Record the WWPN for each port listed.

For more information about using the SANsurfer FC HBA CLI client, see the *SANsurfer FC 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: **`/usr/sbin/hbanyware/hbacmd ListHBAs`**
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 FCP 3

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 FCP service is licensed and started.
- For active/active storage system configurations, verify that the storage system is running a supported cfmode by running the `fcp 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 33.

  For standby and mixed cfmode, special Fibre Channel switch zoning is required. See "Zoning with cfmode standby and mixed" on page 42.
- 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 FCP igroup type as `linux`, and be sure a LUN is mapped as LUN 0.
- **Optional:** Enable ALUA, if you have not already done so.
- 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.

Migrating a storage system to single_image mode 4

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 FCP 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 on the Linux host.

Topics in this chapter

This chapter includes the following topics:

- "Understanding cfmode" on page 34
- "Migrating to single_image mode" on page 36

Understanding cfmode

What cfmode is

Cluster failover mode (cfmode) is a storage system FCP 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 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 FCP 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 FCP 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 FCP switch.

Reasons for using single_image mode

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

- To increase the number of paths to the 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 FCP 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 FCP-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 38.

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: **`lun config_check -S`** **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: **`priv set advanced`**
4	Stop the FCP service on both systems. On each system, enter the following command: **`fcp stop`**
5	Set the cfmode to single _image. On each system console, enter the following command: **`fcp set cfmode single_image`** If you see messages indicating mapping conflicts, then return to Step 2 to resolve the mapping conflicts.
6	Start the FCP service on each system in the active/active configuration. On each system, enter the following command: **`fcp start`**
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: **`fcp nodename`**

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: **`fcp config`** For more information about zoning, see "Validating zoning and connectivity" on page 39.
9	Return to administrative privileges. On each storage system, enter the following command: **`priv set admin`**

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 FCP HBA drivers and mount file systems. Verify that all LUNs were re-discovered by running the sanlun utility.

Configuring FCP cabling and switches 5

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 42
- "Direct-connected configurations" on page 44

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 cfmode standby and mixed

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:

```
fcp show cfmode
```

For information on changing to single_image cfmode, see Chapter 4, "Migrating a storage system to single_image mode," on page 33.

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.

Accessing LUNs from a Red Hat Enterprise Linux host

6

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 75.

If you are using the dm-multipath support, the method you use to access LUNs is different from when you are not using multipathing. This chapter has two main sections, one for accessing LUNs with dm-multipath, 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.

Veritas Storage Foundation

If you are using Veritas Storage Foundation, use the Veritas Multipathing 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 `vxdmp restore` daemon should be configured with an interval of 60.

To do this, run the following command:

```
vxdmptune dmp_restore_interval 60
```

On reboot this value will come into effect and will remain persistent.

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

```
vxdmpadm stop restore
vxdmpadm start restore interval=60
```

If you configure the value by running these commands, the value will not be persistent across reboots.

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

```
vxdmpadm stat restored
```

Topics in this chapter

This chapter includes the following topics:

- "Loading the HBA driver on the host" on page 47
- "Configuring the dm-multipath" on page 49
- "Configuring the multipath service to start automatically after reboot" on page 53
- "Accessing LUNs with dm-multipath support" on page 54
- "Accessing LUNs without dm-multipath support" on page 64
- "Stopping the dm-multipath" on page 63
- "Unloading the HBA driver and shutting down" on page 71

Loading the HBA driver on the host

Loading the HBA driver

Once you have at least one LUN mapped to the Linux host, you can load the HBA driver. When the HBA driver is started, it scans the storage systems and discovers all mapped LUNs. If you configured dm-multipath 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 HBA driver.

Starting FCP

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

Be sure you have configured multipathing before starting FCP.

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

Step	Action

1	Enter the following command on the Linux host:
	`modprobe -v qla2300`
	`modprobe -v qla2400`

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

Step	Action
1	Enter the following command on the Linux host: `modprobe -v lpfc`

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 configured dm-multipath 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 FCP

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

Be sure you have configured multipathing before starting FCP.

Loading the QLogic bundled 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	`modprobe -v qla2300` `modprobe -v qla2400`
	You are using Red Hat Enterprise Linux 5	`modprobe -v qla2xxx`

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: `modprobe -v lpfc`

Configuring the dm-multipath

Note

Skip this section if you are not configuring dm-multipath support or if you are using Veritas Storage Foundation.

Editing the configuration file

To configure the dm-multipath, perform 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 in "In Red Hat Enterprise Linux 5" on page 51 or "In Red Hat Enterprise Linux 4" on page 52, depending on your version of the operating system.
	The file does not exist	Create a new file with the entries as given "In Red Hat Enterprise Linux 5" on page 51 or "In Red Hat Enterprise Linux 4" on page 52, depending on your version of the operating system.
2	Save the changes.	

Step	Action

3	Edit the /etc/multipath.conf file to exclude (blacklist) all devices that do not correspond to any NetApp 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.
	
	a. 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 /dev/sda is a local SCSI drive. To obtain its WWID, run the following command:
	
	`scsi_id -gus /block/sda`
	
	The output looks like as given below:
	
	SIBM-ESXSMAW3073NC_FDAR9P66067WJ
	
	Add this entry for blacklisting.
	
	b. Blacklist other devices using the devnode method by adding the following:
	
	`devnode "^hd[a-z]"`
	
	`devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"`
	
	`devnode "^cciss!c[0-9]d[0-9]*"`
	
	Example: On a Red Hat Enterprise Linux 4 host:
	
	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
	
	}

Step	Action
4	If you are using ALUA, specify the ALUA callout program, that is, the /sbin/mpath_prio_alua program, instead of /sbin/mpath_prio_ontap. 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.

In Red Hat Enterprise Linux 5:

```
defaults
{
   user_friendly_names   yes
}
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
      }
}
```

Note

directio pathchecker is available from Red Hat Enterprise Linux 5 Update 1 alone, and not with Red Hat Enterprise Linux 5.0. For Red Hat Enterprise Linux 5.0, use *readsector0 pathchecker* itself.

In Red Hat Enterprise Linux 4:

```
defaults
{
   user_friendly_names   yes
}
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
     }
}
```

Note

There is a bug if *user_friendly_names* is set to *yes* in the multipath.conf file, upto Red Hat Enterprise Linux 4 Update 4. This has been fixed in Red Hat Enterprise Linux 4 Update 5 and later. For more information, see the burt 196513.

For more information about tuning the multipath configuration, see "Tuning the dm-multipath configuration" on page 62.

Starting the multipath service

Note

Skip this section if you are not configuring dm-multipath support or if you are using Veritas Storage Foundation.

Once you have loaded the HBA drivers, start the multipath service on the Linux host.

Enter the following command:

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

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.

Add the multipath service to the boot sequence by entering the following commands on the Linux console:

```
chkconfig --add multipathd
chkconfig multipathd on
```

Accessing LUNs with dm-multipath support

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 55.
2. Identify and view the list of LUNs. For more information, see "Identifying and viewing a list of LUNs" on page 58.
3. Create and mount the file system. For more information, see "Creating and mounting a file system on a multipath device" on page 60.
4. Tune the multipath configuration if required. For more information, see "Tuning the dm-multipath configuration" on page 62.

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

Verify that the multipath -d 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/sd*x* 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/sd*x* devices are not. After a reboot, or a reload of the HBA driver, you might find that different /dev/sd*x* 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 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 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 both native device-mapper multipath and Veritas Storage Foundation dynamic multipathing solutions, 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   (360a98000433468525634447 17a513571)     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
```

Example: The following output shows when Veritas multipathing is configured.

```
[root@lnx199-115 ~]# sanlun lun show -p
f3020-92:/vol/vol1/115-lun2 (LUN 1)             Lun state: GOOD
Lun Size:      2g (2147483648)   Filer_CF_State: Cluster Enabled
Protocol: FCP           Filer Partner: f3020-93
DMP Nodename: FAS30201_1
Multipath-provider: VERITAS
----------- ------ ---------- ------- ---------- --------------- ---------------
Veritas     sanlun      Filer                     Primary   Partner    Veritas
path          path       Path   /dev/    Host     Filer     Filer        IO
state        state       type    node    HBA      port      port        Policy
----------- ------ ---------- ------- ------------ --------------- ---------------
ENABLED(A)      UP    PRIMARY     sdd    host0    1b        --       Round-Robin
ENABLED         UP  SECONDARY     sdb    host0    --        1a       Round-Robin
ENABLED         UP  SECONDARY     sdh    host1    --        1b       Round-Robin
ENABLED(A)      UP    PRIMARY     sdf    host1    1a        --       Round-Robin
```

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: **`mkdir`** ***`mount_point`*** *mount_point* 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 /etc/fstab as follows, separating each entry with Tab (except the final two zeros, which are separated by a space): ***device*** ***mount_point*** ***type*** **defaults** **0 0** *device* is the name of the device in the /dev/mpath/ directory. *mount_point* is the mountpoint you created for the file system. *type* is the file system type, such as `ext2 or ext3`. **Example:** `/dev/mapper/360a98000433468626734437447 4d4f75 /mnt/ss1/lun_0 ext3 defaults 0 0`
4	To mount the file system, enter the following command: **mount *mountpoint*** **Example:** `mount /mnt/ss1/lun_0`
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 may 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 the dm-multipath” on page 49 for details on blacklisting the local hard drives and other resources that should not be included in the multipathing configuration.

Stopping the dm-multipath

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 FCP on the host, complete the following steps.

Step	Action
1	Unmount all file systems on the FCP 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 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, 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 cfmode standby and mixed" on page 42. 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 65.
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 66.
3. Create a file system on the partition. For more information, see "Creating a file system" on page 68.
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 69.

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 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 FCP Linux Host Utilities. It 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 FCP Linux Host Utilities.

Step	Action

1	Enter the following command at the host's shell prompt:
	`# sanlun lun show all`
	Result: The `sanlun` command displays information about FCP 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`
	`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`

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.`*** *device* 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 `fdisk` 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 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.

Using mkfs to create a file system

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

Step	Action
1	Enter the following command: **`mkfs -t ext3 device`** *device* is the partition. For example, /dev/sda1 is the first partition on /dev/sda.

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 58 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 as follows.

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

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:** `mkdir /mnt/filer0_lun0`
2	Open the file /etc/fstab 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):
	`LABEL=`*label* *mount_point* `auto defaults 0 0`
	label is the file system label you created.
	mount_point is the mountpoint you created for the file system.
	Example:
	`LABEL=lun0_label /mnt/filer0_lun0 auto defaults 0 0`

4	Save your changes to /etc/fstab and exit the editor.
5	Enter **`mount LABEL=`*`label`*** to mount the file system. The file system will be 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 HBA driver and shutting down

Before unloading the HBA driver

Before you unload the HBA driver, ensure that the 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 HBA driver is in use by using the `lsmod` command.

If you are using Veritas Storage Foundation, before unloading the QLogic driver, complete the following steps.

Step	Action
1	Disable the vxconfigd by using the following command: **`vxdctl disable`** **Note** You can verify it using the `vxdctl mode` command.
2	Stop the vxesd, vxrelocate and vxnotify processess.

Unloading the HBA driver

To unload the HBA driver, complete the following steps.

Step	Action
1	**QLogic driver:** Enter the appropriate command for your HBA model: **`modprobe -r qla2400`** **`modprobe -r qla2300`** **Emulex driver:** Enter the following command: **`modprobe -r lpfc`**

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: **`modprobe -r qla2400`** **`modprobe -r qla2300`** If you are using Red Hat Enterprise Linux 5: **`modprobe -r qla2xxx`**
	Bundled Emulex driver: Enter the following command: **`modprobe -r lpfc`**

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`

Accessing LUNs from a SUSE Linux Enterprise Server host

7

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 45.

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

Note

You must use dm-multipath 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 HBA driver on the host" on page 76
- "Configuring the dm-multipath" on page 77
- "Starting the multipath service" on page 81
- "Accessing LUNs with dm-multipath support" on page 82
- "Stopping dm-multipath" on page 91
- "Accessing LUNs without dm-multipath support" on page 92

Loading the HBA driver on the host

Before you load the driver

Once you have at least one LUN mapped to the Linux host, you can load the HBA driver. When the driver is loaded, it scans the storage systems and discover all mapped LUNs. If you configured the dm-multipath 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 HBA driver.

Starting FCP

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

Be sure you have configured multipathing before starting FCP.

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

Step	Action

1	Enter the following command on the Linux host:
	`modprobe -v qla2300`
	`modprobe -v qla2400`

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

Step	Action
1	Enter the following command on the Linux host: `modprobe -v lpfc`

Configuring the dm-multipath

Note

Skip this section if you are not configuring dm-multipath support.

Editing the configuration file

To configure the dm-multipath, perform 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 "In SUSE Linux Enterprise Server 10:" on page 79 or "In SUSE Linux Enterprise Server 9:" on page 79, 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 "In SUSE Linux Enterprise Server 10:" on page 79 or "In SUSE Linux Enterprise Server 9:" on page 79, depending on your version of the operating system.
2	Save the changes.	

Step	Action

3	Edit the /etc/multipath.conf file to exclude (blacklist) all devices that do not correspond to any NetApp 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.
	
	a. 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 /dev/sda is a local SCSI drive. To obtain its WWID, run the following command:
	
	`scsi_id -gus /block/sda`
	
	The output looks like as given below:
	
	SIBM-ESXSMAW3073NC_FDAR9P66067WJ
	
	Add this entry for blacklisting.
	
	b. Blacklist other devices using the devnode method by adding the following:
	
	`devnode "^hd[a-z]"`
	
	`devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"`
	
	`devnode "^cciss!c[0-9]d[0-9]*"`
	
	Example: On a SUSE Linux Enterprise Server 9 host:
	
	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
	
	}

Step	Action
4	Save the changes.

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

In SUSE Linux Enterprise Server 10:

```
defaults
{
   user_friendly_names no
}
blacklist
{
   wwid SIBM-ESXSMAW3073NC_FDAR9P66067WJ
   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"
   path_grouping_policy   group_by_prio
   harware_handler        "0"
   rr_weight              uniform
   getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
   prio_callout          "/sbin/mpath_prio_ontap /dev/%n"
   features              "1 queue_if_no_path"
   path_checker          directio
   failback               immediate
   rr_min_io             128
}
}
```

In SUSE Linux Enterprise Server 9:

```
defaults
{
   user_friendly_names no
}
devnode_blacklist
{
  wwid       SIBM-ESXSMAW3073NC_FDAR9P66067WJ
```

```
  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"
  path_grouping_policy group_by_prio
  rr_weight            uniform
  hardware_handler     "0"
  getuid_callout       "/sbin/scsi_id -g -u -s /block/%n"
  prio_callout         "/sbin/mpath_prio_ontap /dev/%n"
  features             "1 queue_if_no_path"
  path_checker         readsector0
  failback             immediate
  rr_min_io            128
}
}
```

For more information about tuning the multipath configuration, see “Tuning the dm-multipath configuration” on page 90.

Starting the multipath service

Note

Skip this section if you are not configuring dm-multipath support.

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

Step	Action
1	Start the multipath service by entering the following command: `/etc/init.d/boot.multipath start`
2	Start the multipath daemon by entering the following command: `/etc/init.d/multipathd start`

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:
	`insserv boot.multipath multipathd`
	`chkconfig --add multipathd`
	`chkconfig multipathd on`

Accessing LUNs with dm-multipath support

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, dm-multipath configuration" on page 83.
2. Identify and view the list of LUNs. For more information, see "Identifying and viewing a list of LUNs" on page 86.
3. Create and mount the file system. For more information, see "Creating and mounting a file system on a LUN" on page 88.
4. Tune the dm-multipath configuration if required. For more information, see "Tuning the dm-multipath configuration" on page 90.

Verifying the HBA setting, 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 the 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 that the `-d` (“dry run”) parameter prevents the command from actually updating the multipath maps.

Verify that the multipath -d 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/sd*x* 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/sd*x* devices are not. After a reboot, or a reload of the HBA driver, you might find that different /dev/sd*x* 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 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
```

```
[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 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 mount_point`** *mount_point* 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 /etc/fstab as follows, separating each entry with a tab (except the final two zeros, which are separated by a space):
	device ***mount_point*** ***type*** **defaults** **0 0**
	device is the name in the /dev/mapper/ directory.
	mount_point is the mountpoint you created for the file system.
	type is the file system type, such as `ext2`, `ext3`, or `reiserfs`.
	Example:
	`/dev/mapper/360a9800043346536684a337a7553424f`
	`/mnt/ss1/lun_0 ext3 defaults 0 0`

4	Enter the following command to mount the file system: **`mount -a`**
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 the dm-multipath” on page 77 for details on blacklisting the local hard drives and other resources that should not be included in the multipathing configuration.

Stopping dm-multipath

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 FCP on the host, complete the following steps.

Step	Action
1	Unmount all file systems on the FCP 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 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, 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 cfmode standby and mixed" on page 42. 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 93.
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 95.
3. Create a file system on the partition. For more information, see "Creating a file system" on page 96.
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 "Labelling and mounting the file system" on page 97.

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 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 FCP Linux Host Utilities. It displays the host device names and the storage system LUNs to which they are mapped.

Using 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 FCP 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`*** *fcp_device* is the link that maps the FCP 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: **`mkfs -V -t filesystem_type partition`** *partition* 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. *filesystem_type* 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. `mkfs -V -t ext3 /dev/disk/by-id/360a98000433464624f4a34732f785a47-0p1`

Labelling 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 93 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`** *device* is the partition where you created the file system. *label* is a string (maximum 16 characters) used to identify the file system. **Example:** `/sbin/e2label /dev/sda1 lun0_label`

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:** `/dev/disk/by-id/360a98000433464624f4a34732f785a47-0p1 /mnt_pt ext3 defaults 0 0` *mnt_pt* 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 HBA driver and shutting down

Before unloading the HBA driver

Before you unload the HBA driver, ensure that the 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 HBA driver is in use by using the `lsmod` command.

Step	Action
1	Disable the vxconfigd by using the following command: **`vxdctl disable`** **Note** You can verify it using the `vxdctl mode` command.
2	Stop the vxesd, vxrelocate and vxnotify processess.

Unloading the HBA driver

To unload the HBA driver, complete the following steps.

Step	Action
1	**QLogic driver:** Enter the appropriate command for your HBA model: **`modprobe -r qla2400`** **`modprobe -r qla2300`** **Emulex driver:** Enter the following command: **`modprobe -r lpfc`**

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`

Configuring SAN Boot on Red Hat Enterprise Linux host

8

Note

For information on 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 NetApp 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 native host" on page 102
- "Configuring SAN boot for Veritas Storage Foundation" on page 108

Configuring SAN boot on native host

Tasks for configuring SAN boot

To configure SAN boot, perform the following tasks.

1. Install the operating system on the NetApp LUN.
2. Install the Host Utilities.
3. Install and configure the external HBA driver or the bundled HBA driver.
4. Configure dm-multipath as per NetApp recommendations.

Installing the operating system

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

Step	Action
1	On the NetApp storage system, choose a cfmode setting.
2	Create a LUN on the storage system which will act 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 on how to enable the HBA BIOS, see your HBA vendor-specific documentation. **Note** NetApp recommends enabling 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 NetApp LUN. **Attention** Specify *linux mpath* in the anaconda installer, during the operating system installation. When *linux mpath* 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 FCP Linux Host Utilities 4.1. For information on installing the FCP Linux Host Utilities 4.1, see "Installing the FCP Linux Host Utilities software" on page 16.

Installing and configuring external HBA driver

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

Step	Action
1	Download and install the HBA driver and applications. For more information, see "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 on setting the QLogic HBA parameters, see "QLogic HBA parameters" on page 24. To set the HBA driver parameters, perform the following steps. **a.** Edit the /etc/modprobe.conf. **b.** Append the options `qla2xxx statement` with settings in the modprobe.conf file on a single line; only one option `qla2xxx statement` is allowed. **`ql2xsuspendcount=6 MaxRetriesPerPath=10`** **`MaxRetriesPerIo=10 ql2xfailover=0`** **c.** Save the file and exit. Run the `qlinstall -br -in qla2xxx` command. Reboot the host for loading the above external driver. **Emulex:** **Attention** Run the `lpfc-install` script to install the external Emulex driver on the SAN booted host. Reboot the host for loading the above external driver. No special settings are required for supported Emulex HBAs. The default settings work fine in the case of Emulex.

Installing and configuring 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:** The SANsurfer CLI for Linux should be downloaded and installed from the QLogic Web site. **For Emulex:** The Emulex Linux application kit should be downloaded and installed from the Emulex Web site. Ensure that the application kit version matches the bundled Emulex driver version.
2	Configure the bundled HBA parameter. To configure the bundled HBA parameters, complete the following steps. **a.** Run the following command: `scli` **b.** Set the following HBA parameters: ❖ Port Down Retry Count to 30 ❖ Link Down Timeout to 20 **c.** Configure the QLogic HBA port topology to Point to Point Only while connecting to a NetApp 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. **d.** 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. **e.** Select Commit Changes for this HBA so that the changes are applied to the HBA ports. **f.** Repeat the settings for each HBA port listed.

Step	Action
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 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 Netwok web site at rhn.redhat.com/ for SAN boot support on Red Hat Enterprise Linux 5 Update 1. For more infromation about the errata package, see the Notes section of the NetApp® Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability.

Configure the dm-multipath. For more information on configuring the dm-multipath, see "Configuring the dm-multipath" on page 49.

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 NetApp storage systems, perform the following steps.

Step	Action
1	On the NetApp 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 on 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.

Step	Action
7	**Note** Before installing the OS see the section on rootability in the *Veritas™ Volume Manager Administrator's Guide* for Linux that is shipped along with the software for partitioning information. **Attention** Red Hat Enterprise Linux 4 Update 4 does not have inbox drivers for 4Gb QLogic cards, therefore use the driver kit provided by QLogic Website. 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 will boot from the SAN boot LUN on the storage system through a non-proxy path.
9	Install FCP Linux Host Utilities 4.0. For information on installing the FCP Linux Host Utilities 4.0, see "Installing the FCP Linux Host Utilities software" on page 16.

Step	Action
10	Install the supported version of the drivers. For more information about the installation, see "Installing HBAs" on page 21. **QLogic:** **a.** Open the modprobe.conf file. **b.** Append the `options qla2xxx` statement with settings in the modprobe.conf file on a single line; only one `options qla2xxx` statement is allowed: **`ql2xsuspendcount=6 MaxRetriesPerPath=10 MaxRetriesPerIo=10 ql2xfailover=0`** **c.** Save the file and exit. **d.** Set the QLogic parameters to the recommended values. For more information, see "QLogic HBA parameters" on page 24. **e.** Run the following command to reflect the changes in the initrd image. **`qlinstall -up`** This will ensure that the initrd image is updated with the new driver version without unloading the driver. **f.** Reboot the host so that the new driver parameters come into effect. **Emulex:** **a.** Run the following command to install the Emulex driver. **`lpfc-install`** **Note** The `lpfc-install` command will only update mkinitrd with the requisite driver version and will not unload the driver. **b.** Reboot the host to load the requisite driver.

Step	Action
11	Verify the HBA settings. For more information, see "Verifying the HBA settings" on page 55.
12	Install Veritas Storage Foundation and respective ASL and APM.
13	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 *Veritas™ Volume Manager Administrator's Guide* for Linux that is shipped along with the software.
14	Reboot the host after encapsulation.
15	Verify the encapsulation by issuing the following command: **`vxprint`** It will display the rootvol and swapvol under the corresponding disk group.
16	To make the `vxdmpadm restore` daemon interval value persistent in a SAN booted host, complete the following steps: **a.** Edit the /etc/vx/vxvm-startup file. **b.** Add the following entries after the `xconfigd_opts=$opts` line: **`/usr/sbin/vxconfigd -k -m enable`** **`/usr/sbin/vxdmpadm stop restore`** **`/usr/sbin/vxdmpadm start restore interval=60`**
17	Reboot the host.

Step	Action
18	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.

Veritas Array Support Library and Array Policy Module

9

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). 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 114
- "Installing the ASL and APM" on page 115
- "Removing the ASL and APM" on page 119
- "Obtaining multipathing information" on page 120
- "Understanding ASL error message types" on page 123

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.
- Determining the ASL and APM version if it is currently installed
- Obtaining the ASL and APM
- Installing a fresh version of the ASL and APM or upgrading an existing ASL and APM version

Determining the ASL and APM version

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

Example: The following is an example of output from the `vxddladm listversion` command.

```
[root@lnx162 ~]# vxddladm listversion
LIB_NAME                        ASL_VERSION      Min. VXVM version
==================================================================
libvxnetapp.so                    vm-5.0-rev-1            5.0
libvxhitachi.so                   vm-5.0-rev-1            5.0
libvxcscovrts.so                  vm-5.0-rev-1            5.0
libvxibmds6k.so                   vm-5.0-rev-1            5.0
libvxhds9980.so                   vm-5.0-rev-1            5.0
libvxibmds4k.so                   vm-5.0-rev-2            5.0
libvxveritas.so                   vm-5.0-rev-1            5.0
libvxshark.so                     vm-5.0-rev-1            5.0
libvxhdsusp.so                    vm-5.0-rev-2            5.0
libvxibmds8k.so                   vm-5.0-rev-1            5.0
libvxhpalua.so                    vm-5.0-rev-1            5.0
libvxxp1281024.so                 vm-5.0-rev-1            5.0
libvxCLARiiON.so                  vm-5.0-rev-1            5.0
libvxhdsalua.so                   vm-5.0-rev-1            5.0
libvxemc.so                       vm-5.0-rev-2            5.0
libvxxp256.so                     vm-5.0-rev-1            5.0
libvxhds.so                       vm-5.0-rev-1            5.0
libvxpurple.so                    vm-5.0-rev-1            5.0
```

```
libvxxp12k.so                    vm-5.0-rev-2            5.0
```

If you currently have the APM installed, use the Veritas `vxdmpadm listapm all` command to determine its version.

Example: The following shows output from the `vxdmpadm listapm all` command.

```
[root@lnx162 ~]# vxdmpadm listapm all
Filename         APM Name     APM Version  Array Types      State
=================================================================
dmpap.ko         dmpap          1          A/P              Active
dmpap.ko         dmpap          1          A/P-C            Active
dmphpalua.ko     dmphpalua      1          A/A-A-HP         Not-Active
dmphdsalua.ko    dmphdsalua     1          A/A-A-HDS        Not-Active
dmpjbod.ko       dmpjbod        1          Disk             Active
dmpjbod.ko       dmpjbod        1          APdisk           Active
dmpaaa.ko        dmpaaa         1          A/A-A            Not-Active
dmpCLARiiON.ko   dmpCLARiiON    1          CLR-A/P          Not-Active
dmpCLARiiON.ko   dmpCLARiiON    1          CLR-A/PF         Not-Active
dmpaa.ko         dmpaa          1          A/A              Active
dmpnetapp.ko     dmpnetapp      1          A/A-NETAPP       Active
dmpnetapp.ko     dmpnetapp      1          A/P-C-NETAPP     Active
dmpnetapp.ko     dmpnetapp      1          A/P-NETAPP       Active
dmpapg.ko        dmpapg         1          A/PG             Not-Active
dmpapg.ko        dmpapg         1          A/PG-C           Not-Active
dmpapf.ko        dmpapf         1          A/PF-VERITAS     Not-Active
dmpapf.ko        dmpapf         1          A/PF-T3PLUS      Not-Active
```

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

Step	Action
1	Log in to the VxVM system as the root user.

Step	Action
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. You should have your LUNs set up before you install the ASL and APM.

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 `vxdmpadm listenclosure all` command. The output shows the model name of the storage device if you are using enclosure-based naming with VxVM.
	`vxdmpadm listenclosure all`
	Note
	After you verify the installation, you can use the `vxdctl enable` command to restart the `vxconfigd` daemon. You only use this command when your host is connected to NetApp storage systems.
	Example: The `vxdmpadm listenclosure all` command shows the Enclosure Types as FAS3070 in this example.
	`[root@lnx162 ~]# vxdmpadm listenclosure all`
	`ENCLR_NAME  ENCLR_TYPE  ENCLR_SNO  STATUS  ARRAY_TYPE`
	`==========`
	`Disk  Disk  DISKS  CONNECTED  Disk`
	`FAS30700 FAS3070 3079851  CONNECTED  A/P-C-NETAPP`

Step	Action

5	If your host is not connected to storage, use the following command:
	`vxddladm listsupport all`
	Example: The following is a sample of the output you see when you enter the `vxddladm listsupport all` command.
	`[root@lnx162 ~]# vxddladm listsupport all`
	`LIBNAME VID`
	`==========`
	`libvxnetapp.so NETAPP`
	`libvxhitachi.so HITACHI`
	`libvxcscovrts.so CSCOVRTS`
	`libvxibmds6k.so IBM`
	`libvxhds9980.so HITACHI`
	`libvxibmds4k.so IBM`
	`libvxveritas.so VERITAS`
	`libvxshark.so IBM`
	`libvxhdsusp.so HITACHI`
	`libvxibmds8k.so IBM`
	`libvxhpalua.so HP, COMPAQ`
	`libvxxp1281024.so HP`
	`libvxCLARiiON.so DGC`
	`libvxhdsalua.so HITACHI`
	`libvxemc.so EMC`
	`libvxxp256.so HP`
	`libvxhds.so HITACHI`
	`libvxpurple.so SUN`
	`libvxxp12k.so HP`

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-FCP-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 NetApp LUNs. **Note** In a Storage Foundation RAC cluster, you must also stop clustering on a node before you remove the ASL and APM.

3	◆ Enter the following command to remove the ASL package:
	`rpm -e` ***`asl_rpm_name`***
	Examples:
	The following command removes a previous version of the ASL.
	`rpm -e VRTSNTAPasl-5.0-1.0_RHEL4`
	◆ Enter the following command to remove the APM package:
	`rpm -e` ***`apm_rpm_name`***
	Example:
	The following command removes a previous version of the APM.
	`rpm -e VRTSNTAPapm-5.0-1.0_RHEL4`

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:
	`vxdisk list`
	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.
	`[root@lnx162 ~]# vxdisk list`
	`DEVICE TYPE DISK GROUP STATUS`
	`FAS30700_0 auto:cdsdisk - - online`
	`FAS30700_2 auto:cdsdisk FAS30700_4 dg2 online shared`
	`FAS30700_3 auto:cdsdisk - - online`
	`FAS30700_4 auto:cdsdisk - - online`
	`FAS30700_5 auto:cdsdisk FAS30700_0 dg1 online shared`
	`FAS30700_10 auto:cdsdisk FAS30700_8 dg1 online shared`
	`FAS30700_11 auto:cdsdisk FAS30700_7 dg2 online shared`
	`sda auto:none - - online invalid`

Step	Action

2	On the host console, enter the following command to display path information for the device you want:
	`vxdmpadm getsubpaths dmpnodename=`*`device`*
	device is the name listed under the output of the `vxdisk list` command.
	Example: The following is sample output.
	`[root@lnx162 ~]# vxdmpadm getsubpaths dmpnodename=FAS30700_11`
	`NAME STATE[A]  PATH-TYPE[M] CTLR-NAME ENCLR-TYPE ENCLR-NAME ATTRS`
	`==========`
	`sdab ENABLED(A) PRIMARY  c3  FAS3070 FAS30700 -`
	`sdg  ENABLED(A) PRIMARY  c2  FAS3070 FAS30700 -`
	`sdn  ENABLED  SECONDARY  c2  FAS3070 FAS30700 -`
	`sdu  ENABLED  SECONDARY  c3  FAS3070 FAS30700 -`

3	To obtain path information for a host HBA, enter the following command:
	`vxdmpadm getsubpaths ctlr=`*`controller_name`*
	controller_name is the controller displayed under “CTRL-NAME” in the output of the `vxdmpadm getsubpaths dmpnodename` 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.
	`[root@lnx162 ~]# vxdmpadm getsubpaths ctlr=c2`
	`NAME  STATE[A]  PATH-TYPE[M] DMPNODENAME  ENCLR-TYPE ENCLR-NAME  ATTRS`
	`==========`
	`sdb  ENABLED  SECONDARY  FAS30700_5 FAS3070 FAS30700 -`
	`sdc  ENABLED  SECONDARY  FAS30700_2 FAS3070 FAS30700 -`
	`sdd  ENABLED  SECONDARY  FAS30700_0 FAS3070 FAS30700 -`
	`sde  ENABLED  SECONDARY  FAS30700_3 FAS3070 FAS30700 -`
	`sdf  ENABLED  SECONDARY  FAS30700_4 FAS3070 FAS30700 -`
	`sdg  ENABLED(A) PRIMARY  FAS30700_11 FAS3070 FAS30700 -`
	`sdh  ENABLED(A) PRIMARY  FAS30700_10 FAS3070 FAS30700 -`
	`sdi  ENABLED(A) PRIMARY  FAS30700_5 FAS3070 FAS30700 -`
	`sdj  ENABLED(A) PRIMARY  FAS30700_2 FAS3070 FAS30700 -`
	`sdk  ENABLED(A) PRIMARY  FAS30700_0 FAS3070 FAS30700 -`
	`sdl  ENABLED(A) PRIMARY  FAS30700_3 FAS3070 FAS30700 -`
	`sdm  ENABLED(A) PRIMARY  FAS30700_4 FAS3070 FAS30700 -`
	`sdn  ENABLED  SECONDARY  FAS30700_11 FAS3070 FAS30700 -`
	`sdo  ENABLED  SECONDARY  FAS30700_10 FAS3070 FAS30700 -`

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:
	`sanlun lun show -p`
	Result: `sanlun` displays path information for each LUN.
	Example: The following is sample output of the `sanlun lun show -p` command:
	`f3070-199-205:/vol/vol1/l1 (LUN 2)  Lun state: GOOD`
	`Lun Size:  100m (104857600)  Filer_CF_State: Cluster Enabled`
	`Protocol: FCP  Filer Partner: f3070-199-206`
	`DMP Nodename: FAS30700_0`
	`Multipath-provider: VERITAS`
	`----------- ------ ---------- ------- ------------ --------------- ----------`
	`Veritas  sanlun  Filer  Primary  Partner  Veritas`
	`path  path  Path  /dev/  Host  Filer  Filer  IO`
	`state  state  type  node  HBA  port  port  Policy`
	`----------- ------ ---------- ------- ---------- --------------- ------------`
	`ENABLED  UP  SECONDARY  sdy  host3  --  0a  Round-Robin`
	`ENABLED(A)  UP  PRIMARY  sdr  host3  0a  --  Round-Robin`
	`ENABLED(A)  UP  PRIMARY  sdk  host2  0c  --  Round-Robin`
	`ENABLED  UP  SECONDARY  sdd  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 A

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 FCP 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 fails.

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

Linux FCP hosts with QLogic initiators fail to detect LUNs

Problem: Linux FCP 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 FCP host, ensure the following settings are done for port topology and speed using the SANsurfer scli package as follows:

1. Hard set the QLogic HBA port topology to Point to Point Only while connecting to a NetApp filer 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:

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

2. Hard set the 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:

   - Run the `/opt/QLogic_Corporation/SANsurferCLI/scli`
   - 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

FCP 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 HBA driver version you are running is supported. You can check the most up-to-date system requirements in the NetApp FCP 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 properly connected.
Zoning	Verify that the zoning is correct. For information on zoning see "Switched fabric configurations" on page 42.
HBA Setting	Verify the HBA settings. For information on verifying the HBA settings, see "Verifying the HBA settings" on page 55.
FCP service status	Verify that the FCP service is licensed and started on the storage system according to the procedure described in the Data ONTAP *Block Access Management Guide*.
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.

What to do: Make sure 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 FCP 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 on blacklisting devices on a Red Hat Enterprise Linux host, see "Configuring the dm-multipath" on page 49.

For more information about blacklisting devices on a SUSE Linux Enterprise Server host, see "Configuring the dm-multipath" on page 77.

If the FCP 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.

FCP performance is poor

Problem: If you experience poor performance on some or all FCP 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.

What to do: 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
Make sure the package installing the library is installed & loaded
```

What to do: 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.

What to do: 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
```

Troubleshooting — Veritas Storage Foundation

B

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 hastop on all the nodes.
2	Try stopping the fencing driver again.
3	If it does not work, remove the fencdg directory from /etc/vxfendg.
4	In the /etc/vxfenmode file, make the following change: **`vxfen_mode=disabled`**
5	Stop the fencing driver by using the following command: **`/etc/init.d/vxfen stop`**

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 `vxdmp` deamon will store 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:

```
vxdmpadm settune dmp_log_level=4
```