

# FCP Linux® Host Utilities 3.0 for Native OS **Setup Guide**

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

Part number 215-02003\_A0  
August 2006

# Copyright and trademark information

---

## Copyright information

Copyright © 1994–2006 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

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

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

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

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

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

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

## Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, Data ONTAP, FilerView, MultiStore, NearStore, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapVault, SyncMirror, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The Evolution of Storage are trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. FlexClone, FlexVol, NOW, NOW NetApp on the Web, RAID-DP, Smart SAN, SnapCache, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, SnapValidator, vFiler, VFM, and VFM (Virtual File Manager) are trademarks of Network Appliance, Inc. in the United States and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the U.S.A.

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

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

Network Appliance NetCache is certified RealSystem compatible.

# Table of Contents

---

|                  |                                                                       |    |
|------------------|-----------------------------------------------------------------------|----|
|                  | <b>Preface</b> . . . . .                                              | v  |
| <b>Chapter 1</b> | <b>Setup Overview</b> . . . . .                                       | 1  |
|                  | Setup procedure. . . . .                                              | 2  |
|                  | Where to find the information you need . . . . .                      | 4  |
|                  | Understanding the Host Utilities components and requirements. . . . . | 6  |
|                  | Understanding host and storage system communication . . . . .         | 8  |
| <b>Chapter 2</b> | <b>Installing HBAs and software on the host</b> . . . . .             | 11 |
|                  | Installing the Host Utilities . . . . .                               | 12 |
|                  | Removing the previous HBA drivers. . . . .                            | 14 |
|                  | Verifying or installing the required RPMs . . . . .                   | 15 |
|                  | Installing host bus adapters . . . . .                                | 16 |
|                  | Downloading and installing the HBA driver and applications . . . . .  | 17 |
|                  | Setting required HBA and driver parameters. . . . .                   | 19 |
|                  | Recording the WWPNs of the HBAs. . . . .                              | 22 |
| <b>Chapter 3</b> | <b>Configuring the storage system for FCP.</b> . . . . .              | 25 |
| <b>Chapter 4</b> | <b>Migrating a storage system to single_image mode</b> . . . . .      | 27 |
|                  | Understanding cfmode . . . . .                                        | 28 |
|                  | Completing prerequisites . . . . .                                    | 30 |
|                  | Upgrading to single_image mode . . . . .                              | 31 |
| <b>Chapter 5</b> | <b>Configuring FCP cabling and switches</b> . . . . .                 | 37 |
|                  | Switched fabric configurations . . . . .                              | 38 |
|                  | Direct connected configurations . . . . .                             | 39 |

|                       |                                                                      |               |
|-----------------------|----------------------------------------------------------------------|---------------|
| <b>Chapter 6</b>      | <b>Accessing LUNs from a RHEL host . . . . .</b>                     | <b>41</b>     |
|                       | Loading the HBA driver on the host . . . . .                         | 42            |
|                       | Editing the host's /etc/multipath.conf file . . . . .                | 43            |
|                       | Configuring the multipath service to start automatically. . . . .    | 47            |
|                       | Accessing LUNs using dm-multipath . . . . .                          | 48            |
|                       | Viewing a list of LUNs. . . . .                                      | 49            |
|                       | Viewing the dm-multipath configuration and multipath devices . . . . | 51            |
|                       | Creating and mounting a file system on a multipath device . . . . .  | 52            |
|                       | Accessing LUNs as a raw devices . . . . .                            | 54            |
|                       | Tuning the dm-multipath configuration . . . . .                      | 55            |
|                       | Accessing LUNs without dm-multipath . . . . .                        | 56            |
|                       | Creating a partition . . . . .                                       | 58            |
|                       | Creating a file system . . . . .                                     | 60            |
|                       | Labeling and mounting the file system. . . . .                       | 61            |
|                       | Accessing LUNs as raw devices . . . . .                              | 63            |
|                       | Viewing LUN information . . . . .                                    | 64            |
|                       | Stopping or reloading the HBA driver and shutting down . . . . .     | 65            |
| <br><b>Chapter 7</b>  | <br><b>Accessing LUNs from a SLES host. . . . .</b>                  | <br><b>67</b> |
|                       | Loading the HBA driver on the host . . . . .                         | 68            |
|                       | Editing the host's /etc/multipath.conf file . . . . .                | 69            |
|                       | Starting the multipath service. . . . .                              | 73            |
|                       | Configuring the multipath service to start automatically. . . . .    | 74            |
|                       | Accessing LUNs with dm-multipath support. . . . .                    | 75            |
|                       | Verifying the dm-multipath policy and mapping . . . . .              | 76            |
|                       | Viewing a list of LUNs. . . . .                                      | 77            |
|                       | Creating and mounting a file system on a LUN . . . . .               | 79            |
|                       | Accessing LUNs as a raw devices . . . . .                            | 81            |
|                       | Tuning the dm-multipath configuration . . . . .                      | 82            |
|                       | Stopping dm-multipath. . . . .                                       | 83            |
|                       | Accessing LUNs without dm-multipath . . . . .                        | 84            |
|                       | Creating a file system on a LUN . . . . .                            | 88            |
|                       | Accessing LUNs as raw devices . . . . .                              | 91            |
|                       | Stopping or reloading the HBA driver and shutting down . . . . .     | 92            |
| <br><b>Appendix A</b> | <br><b>Troubleshooting . . . . .</b>                                 | <br><b>93</b> |

# Preface

---

## About this guide

This guide describes how to install the FCP Linux® Host Utilities 3.0 for Native OS on a Linux host and how to set up the host to access LUNs on a NetApp storage system. It also explains how to configure the storage system to work with the host.

The FCP Linux Host Utilities 3.0 for Native OS is the follow-on release to the FCP Linux Attach Kit 1.0.

See the *NetApp FCP Host Compatibility Matrix*, available on the NetApp on the Web™ (NOW) site at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/) for information on the specific Linux versions supported.

## Audience

This guide is for administrators of Red Hat Enterprise Linux or SuSE Linux Enterprise Server hosts and NetApp storage systems.

## Terminology

This guide uses the following terms:

- ◆ NetApp storage products (filers, FAS appliances, and NearStore systems) are all *storage systems*—also sometimes called filers or storage appliances.
- ◆ The term *type* means pressing one or more keys on the keyboard. The term *enter* means pressing one or more keys and then pressing the Enter key.

## 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                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Italic type</i>          | <ul style="list-style-type: none"> <li>◆ Words or characters that require special attention.</li> <li>◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fcstest adaptername</code> command, you enter the characters “fcstest” followed by the actual name of the adapter.</li> <li>◆ Book titles in cross-references.</li> </ul> |
| Monospaced font             | <ul style="list-style-type: none"> <li>◆ Command and daemon names.</li> <li>◆ Information displayed on the system console or other computer monitors.</li> <li>◆ The contents of files.</li> </ul>                                                                                                                                                                                        |
| <b>Bold monospaced font</b> | Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.                                                                                                                                                                                                      |

## Keyboard conventions

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

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

## Special messages

This guide contains special messages that are described as follows:

---

### **Note**

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

---

---

### **Caution**

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

---



**About this chapter**

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

**About Linux versions**

The FCP Linux® Host Utilities support both the Red Hat® Enterprise Linux (RHEL) and SuSE® Linux Enterprise Server (SLES) operating systems. Where the setup procedure differs based on Linux version, this information is called out in this guide. If not otherwise marked, the information in this guide applies to all supported versions of Linux software.

**About Host Bus Adapters**

The FCP Linux® Host Utilities support both QLogic® and Emulex® brand Host Bus Adapters (HBAs). Where the setup procedure differs based on HBA model, this information is called out in this guide. If not otherwise marked, the information in this guide applies to all supported versions HBA models. Do not use different HBA brands on a single Linux host.

**Topics in this chapter**

This chapter includes the following topics:

- ◆ [“Setup procedure”](#) on page 2
- ◆ [“Where to find the information you need”](#) on page 4
- ◆ [“Understanding the Host Utilities components and requirements”](#) on page 6
- ◆ [“Understanding host and storage system communication”](#) on page 8

# Setup procedure

---

## Before you begin

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

- ◆ Linux operating system version
- ◆ Host bus adapter (HBA) model
- ◆ HBA firmware and driver version
- ◆ Fibre Channel switch model and firmware version
- ◆ Storage system and Data ONTAP version

You can find the interoperability matrixes at [https://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](https://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/)

## Setup steps

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Install and configure HBAs and software on the host, including:</p> <ul style="list-style-type: none"><li>◆ Installing the FCP Linux Host Utilities software.</li><li>◆ Removing any existing HBA drivers.</li><li>◆ Installing dm-multipath RPMs.</li><li>◆ Installing the HBAs if not already installed. Do not cable HBAs; remove cables if already installed.</li><li>◆ Downloading and installing the supported HBA driver.</li><li>◆ Setting required HBA parameters.</li><li>◆ Recording the World Wide Port Names (WWPNs) of the HBAs.</li></ul> <p>See Chapter 2, “<a href="#">Installing HBAs and software on the host</a>,” on page 11.</p> |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2    | <p>Configure the storage system, including:</p> <ul style="list-style-type: none"> <li>◆ Licensing and starting the FCP service.</li> <li>◆ Creating LUNs, creating an igroup that contains the host's FCP WWPN, and mapping the LUNs to that igroup. A LUN 0 must be mapped to the host.</li> </ul> <p>See Chapter 3, "<a href="#">Configuring the storage system for FCP</a>," on page 25.</p>                                                                                               |
| 3    | <p>Optionally migrate an active/active storage system configuration to single-image cfmode.</p> <p>See Chapter 4, "<a href="#">Migrating a storage system to single_image mode</a>," on page 27.</p>                                                                                                                                                                                                                                                                                           |
| 4    | <p>Configuring the Fibre Channel switches and cabling, including:</p> <ul style="list-style-type: none"> <li>◆ Verifying the firmware revision of the switches.</li> <li>◆ Cabling the storage system to the switches.</li> <li>◆ Creating zones on the switches.</li> <li>◆ Cabling the host HBAs to the switches or directly to the storage system after a LUN 0 is mapped to the host.</li> </ul> <p>See Chapter 5, "<a href="#">Configuring FCP cabling and switches</a>," on page 37.</p> |
| 5    | <p>Discovering and accessing the LUNs from the host, including:</p> <ul style="list-style-type: none"> <li>◆ Loading the HBA driver modules.</li> <li>◆ Configuring multipathing using dm-multipath.</li> <li>◆ Creating file systems on the LUNs and mounting them, or configuring the LUNs as raw devices</li> </ul> <p>See Chapter 6, "<a href="#">Accessing LUNs from a RHEL host</a>," on page 41 or Chapter 7, "<a href="#">Accessing LUNs from a SLES host</a>," on page 67.</p>        |

---

**Note**

You must alternate between setting up the host and the storage system in the order shown above.

---

## Where to find the information you need

---

### About 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...                           | Go to...                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Known issues, system requirements, and last minute updates | The <i>Release Notes</i> included with these Host Utilities. Also available at <a href="http://now.netapp.com/NOW/knowledge/docs/san/">http://now.netapp.com/NOW/knowledge/docs/san/</a>                                                                                                                                                                                                                                                                                   |
| The latest supported configurations                        | <ul style="list-style-type: none"><li>◆ The appropriate interoperability matrix at <a href="https://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/">https://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/</a></li><li>◆ <i>System Configuration Guide</i> at <a href="https://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/">https://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/</a></li></ul>                                |
| Configuring the storage system                             | <ul style="list-style-type: none"><li>◆ <i>Best Practices for Reliability</i> at <a href="http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm">http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm</a></li><li>◆ <i>Data ONTAP Software Setup Guide</i> at <a href="http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml">http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml</a></li></ul> |
| Supported Fibre Channel SAN topologies                     | <i>FCP Configuration Guide</i> at <a href="http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/">http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/</a>                                                                                                                                                                                                                                                                                              |

| <b>If you need information about...</b>                              | <b>Go to...</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Installing and configuring the HBA in your host                      | Refer to your HBA vendor documentation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Installing and configuring a supported version of SnapDrive for UNIX | The <i>Installation and Administration Guide</i> for that version of SnapDrive for UNIX.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Managing SAN storage on the storage system                           | <ul style="list-style-type: none"> <li>◆ <i>Command Reference</i> at <a href="http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml">http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml</a></li> <li>◆ <i>Data ONTAP Block Access Management Guide</i> at <a href="http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml">http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml</a></li> <li>◆ <i>FilerView Online Help</i></li> <li>◆ <i>Best Practices for Reliability: Ongoing Maintenance</i> at <a href="http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm">http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm</a></li> </ul> |

# Understanding the Host Utilities components and requirements

---

## What the FCP Linux Host Utilities are

The FCP Linux Host Utilities 3.0 for Native OS enable a Linux host to connect to your storage systems. The kit 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

This FCP Linux Host Utilities consist of the following software components:

- ◆ The `sanlun` utility that displays information about LUNs on the storage system that are available to this Linux host.
- ◆ The `qla2xxx_lun_rescan` script that discovers new LUNs when using a QLogic HBA.
- ◆ An updated `mpath_prio_ontap` callout file that replaces the `/sbin/mpath_prio_netapp` file included in the Linux distribution.
- ◆ The `install` and `uninstall` scripts to add and remove the Host Utilities components.
- ◆ The following diagnostic scripts for collecting troubleshooting information:
  - ❖ `brocade_info` - collects information about Brocade switches installed in the network.
  - ❖ `cisco_info` - collects information about Cisco switches installed in the network.
  - ❖ `filer_info` - collects information about the storage system.
  - ❖ `libHBAAPI.so` - support library for `sanlun`.
  - ❖ `linux_info` - collects configuration information about your Linux host.
  - ❖ `mcddata_info` - collects information about McData switches installed in the network.
  - ❖ `san_version` - displays the version of the Host Utilities.
  - ❖ `SHsupport.pm` - support library for `*info` scripts.
  - ❖ `Telnet.pm` - support library for `*info` scripts.
- ◆ The man pages for `sanlun` and the diagnostic scripts.
- ◆ This *Setup Guide*.
- ◆ The *Release Notes*.
- ◆ Quick command reference guides.

## Multipathing

Linux device mapper multipathing (dm-multipath) is supported on NetApp storage systems with Linux hosts running RHEL 4 Update 2 and Update 3, and SLES 9 SP2 and SP3. 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 the RHEL 4 Update 2 and Update 3 and SLES 9 SP2 and SP3 binary kernel RPMs. You may have to install additional userspace RPMs as part of the setup process.

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

## 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 4Gb target HBA, then you must manually change the cfmode.

See Chapter 4, “[Migrating a storage system to single\\_image mode](#),” on page 27 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 Fibre Channel Protocol for SCSI to access data on storage systems that run supported versions of Data ONTAP software.

The storage system is a FCP target device. Each host HBA is an initiator that uses FCP to access logical units of storage (LUNs) on a storage system to store and retrieve data.

On a host, a World Wide Port Name (WWPN) identifies each port on an HBA. The host WWPNs are used as identifiers when creating initiator groups on a storage system. An initiator group permits host access to specific LUNs on a storage system.

## Supported configurations

The following configurations are supported:

- ◆ Fabric-attached storage area network (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.

---

See the *Network Appliance Fibre Channel Configuration Guide* at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/) for more detailed information, including diagrams, on the supported Fibre Channel topologies.



## Configuration overview

This chapter explains how to install the FCP 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 device mapper multipathing (dm-multipath) support to create a highly-available connection between the Linux host and storage system, you also need the device-mapper (RHEL and SLES) and multipath-tools (SLES) or device-mapper-multipath (RHEL) RPMs.

---

### Note

Do not connect the Fibre Channel cables to the HBAs on the Linux host until instructed in “[Connect cables](#)” on page 38. A storage system LUN 0 must be available to the Linux host when the cable is connected for the first time.

---

## Topics in this chapter

This chapter includes the following topics:

- ◆ “[Installing the Host Utilities](#)” on page 12
- ◆ “[Removing the previous HBA drivers](#)” on page 14
- ◆ “[Verifying or installing the required RPMs](#)” on page 15.
- ◆ “[Installing host bus adapters](#)” on page 16
- ◆ “[Downloading and installing the HBA driver and applications](#)” on page 17
- ◆ “[Setting required HBA and driver parameters](#)” on page 19
- ◆ “[Recording the WWPNs of the HBAs](#)” on page 22

## Requirement for root privileges

Root privileges are required to configure initiator settings. Before you configure the initiator, make sure you are logged in with root privileges.

# Installing the Host Utilities

---

## Verifying the correct version

Verify that you have the correct FCP Host Utilities for your version of Linux. See the *NetApp FCP Host Compatibility Matrix* at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/).

## Copying or downloading the Host Utilities software

Copy the FCP Host Utilities file (netapp\_linux\_host\_utils\_3\_0.tar.gz) from CD or download it from the NOW web site at <http://now.netapp.com/NOW/cgi-bin/software> to a working directory on your Linux host.

## Installing the Host Utilities software

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

| Step | Action                                                                                                                                                                                                   |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Remove any previous Linux host attach kit. Change to the directory where the previous version is installed (default is /opt/sanlun/bin) and enter the following command:<br><br><code>./uninstall</code> |
| 2    | Change to the working directory to which you copied the Host Utilities file.                                                                                                                             |
| 3    | Enter the following command to uncompress the file:<br><br><code>gunzip netapp_linux_host_utils_3_0.tar.gz</code>                                                                                        |
| 4    | Enter the following command to extract the files:<br><br><code>tar -xvf netapp_linux_host_utils_3_0.tar</code>                                                                                           |
| 5    | Change to the netapp_linux_host_utils_3_0 directory. By default, this directory is a subdirectory of the working directory in which you extracted the Host Utilities files in the previous step.         |
| 6    | Enter the following command:<br><br><code>./install</code>                                                                                                                                               |

The diagnostic scripts are installed to the `/opt/netapp/santools` directory. Note that this directory is different from the directory used by the previous version of the support kit.

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:

| Step | Action                                                                    |
|------|---------------------------------------------------------------------------|
| 1    | Log in to the Linux host as root.                                         |
| 2    | Stop qlremote by entering the following command:<br><b>pkill qlremote</b> |
| 3    | Change to the directory where the driver installer files are located.     |
| 4    | Run the following command:<br><b>./qlinstall --uninstall</b>              |

### Removing existing Emulex HBA drivers

You need to remove any existing HBA driver modules. To remove the Emulex driver modules:

| Step | Action                                                                |
|------|-----------------------------------------------------------------------|
| 1    | Log in to the Linux host as root.                                     |
| 2    | Change to the directory where the driver installer files are located. |
| 3    | Run the following command:<br><b>./lpfc-install --uninstall</b>       |

You can also remove Emulex version 8.x drivers using the `lpfc-install --uninstall` command in the new driver package you install in [“Downloading and installing the HBA driver and applications”](#) on page 17. Be sure to uninstall the old drivers before installing the new version.

## Verifying or installing the required RPMs

---

**Multipathing RPMs**      To verify that the correct multipathing RPMs are installed on your Linux host, complete the following steps:

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>If you plan to use dm-multipath support, enter the following commands:</p> <pre>rpm -q device-mapper           (both SLES and RHEL) rpm -q multipath-tools         (SLES hosts only) rpm -q device-mapper-multipath (RHEL hosts only)</pre> <p><b>Result:</b> The <code>rpm -q</code> commands return the name and version of the installed RPMs.</p> <p>See the <i>Release Notes</i> for the correct RPM version for your specific version of Linux.</p> |
| 2    | <p>If you do not have the required RPMs, install them using the appropriate method for your version of Linux.</p>                                                                                                                                                                                                                                                                                                                                            |

# Installing host bus adapters

---

## Verify HBA model is supported

Verify that your HBA model is supported. See the *NetApp FCP Host Compatibility Matrix* at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/).

## 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 could lead to driver conflicts.

## Hardware installation

Follow the HBA and host hardware vendors' instructions for installing the HBAs in the Linux host. If the HBAs are already installed, skip this step.

### **Note**

---

Do not cable the HBAs to the switch or storage system yet. If existing HBAs are cabled, remove the Fibre Channel cables from the HBAs at this time.

---

# Downloading and installing the 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. See the *NetApp FCP Host Compatibility Matrix* at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/).

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 and applications package from the HBA vendor's web site.

**QLogic:** For QLogic HBAs, go to [http://support.qlogic.com/support/drivers\\_software.asp](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.

## Installing 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.

---

### Note

Be sure the Fibre Channel cables are not connected to the HBA before installing the driver.

---

A reboot is required after installing the HBA driver package.

**QLogic HBA:** Use the `./qlinstall` program with no options to install the driver package.

**Emulex HBA:** Use the `./lpfc-install` program with no options to install the driver package.

## Installing HBA management software

You should install the HBA vendor's command line or graphical management software. The examples in this guide assume you have installed the SANsurfer FC HBA CLI software (QLogic) or the HBAnyware CLI Client software (Emulex). You can perform these same tasks using the equivalent graphical user interface if you prefer.

**QLogic:** The QLogic management CLI (`scli`) is installed by the driver installation program. You can optionally install the SANsurfer Pro GUI. The SANsurfer Pro version includes software components used by the `sanlun` utility to display HBA information.

**Emulex:** The management software is included in the Emulex Linux applications kit. The kit also contains software components used by the `sanlun` utility to display HBA information. You must reboot the Linux host after installing this package.

## Setting required HBA and driver parameters

---

### For QLogic HBAs

You need to set both HBA and driver parameters to ensure that multipathing and storage system failover works 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    | Start the <code>scli</code> command in interactive mode by entering the following command:<br><br><code>scli</code>                                                                    |
| 2    | From the main menu, select the Configure HBA Settings option.                                                                                                                          |
| 3    | Select the first HBA port listed.                                                                                                                                                      |
| 4    | Select the option for Port Down Retry Count and set the value to <b>30</b> if you are using <code>dm-multipath</code> , or <b>180</b> if you are not using <code>dm-multipath</code> . |
| 5    | Select the option for Link Down Timeout and set the value to <b>20</b> .                                                                                                               |
| 6    | Repeat the settings for each HBA port listed.                                                                                                                                          |
| 7    | Return to the Main Menu and Quit.                                                                                                                                                      |

**QLogic driver parameters:** To set the driver parameters, complete the following steps on the Linux host. Note that you need to edit either `/etc/modprobe.conf` or `/etc/modprobe.conf.local`, 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:<br><br><code>./qlinstall -ul qla2400</code><br><code>./qlinstall -ul qla2300</code> |
| 2    | Open the <code>/etc/modprobe.conf</code> file with a text editor.                                                                                                        |

| Step | Action                                                                                                                                                                                                                                                                                                                                                      |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | Locate the options <code>qla2xxx</code> statement that was added by the QLogic installer. If the options <code>qla2xxx</code> statement is present, modify it as shown in the next step. If there is no options <code>qla2xxx</code> statement in <code>/etc/modprobe.conf</code> , close that file and edit <code>/etc/modprobe.conf.local</code> instead. |
| 4    | Append the options <code>qla2xxx</code> statement with following settings on a single line; only one options <code>qla2xxx</code> statement is allowed.<br><br><code>ql2xsuspendcount=6 MaxRetriesPerPath=10</code><br><code>MaxRetriesPerIo=10 qla2xfailover=0</code>                                                                                      |
| 5    | Save the file and exit.                                                                                                                                                                                                                                                                                                                                     |
| 6    | Reload the QLogic HBA driver using the appropriate command for your HBA model:<br><br><code>./qlinstall -l qla2400</code><br><code>./qlinstall -l qla2300</code>                                                                                                                                                                                            |
| 7    | Enter the appropriate command for your HBA model to rebuild the ramdisk image with the new parameter values:<br><br><code>./qlinstall -br -in qla2400</code><br><code>./qlinstall -br -in qla2300</code>                                                                                                                                                    |
| 8    | Reboot the Linux host using the updated image.                                                                                                                                                                                                                                                                                                              |

## For Emulex HBAs

No special settings are required for supported Emulex HBAs if you are using `dm-multipath`. Use the factory default settings.

If you are not using `dm-multipath`, complete the following steps to set the required parameter.

| Step | Action                                                                                               |
|------|------------------------------------------------------------------------------------------------------|
| 1    | Unload the Emulex HBA driver by entering the following command:<br><br><code>modprobe -r lpfc</code> |
| 2    | Open the <code>/etc/modprobe.conf</code> file with a text editor.                                    |

| Step | Action                                                                                                                                      |
|------|---------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | Add the following statement:<br><b>options lpfc lpfc_nodev_tmo=180</b>                                                                      |
| 4    | Save the file and exit.                                                                                                                     |
| 5    | Reload the Emulex HBA driver:<br><b>modprobe lpfc</b>                                                                                       |
| 6    | Enter the following command to rebuild the ramdisk image with the new parameter value:<br><b>/usr/src/lpfc/lpfc-install --createramdisk</b> |
| 7    | Reboot the Linux host using the updated image.                                                                                              |

# Recording the WWPNs of the HBAs

## About WWPNs

The storage system uses the World Wide Port Name (WWPN) of the HBA ports to identify them. The WWPN is added to an initiator group (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 Fibre Channel Protocol also defines a World Wide Node Name (WWNN). The WWNN is not used when creating igroups on the storage system.

## Obtaining the WWPN for QLogic HBAs

To obtain the World Wide Port names 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:<br><code>scli -i</code> |
| 2    | Record the WWPN for each port listed.                                          |

For more information on using the SANsurfer FC HBA CLI client, see the *SANsurfer FC HBA CLI Application User's Guide*. To download the guide, point your browser to [http://support.qlogic.com/support/drivers\\_software.asp](http://support.qlogic.com/support/drivers_software.asp) and then click your HBA model.

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

## Obtaining the WWPN for Emulex HBAs

To obtain the World Wide Port names of using the Emulex HBAnyware for Linux CLI, complete the following steps.

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

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



## Before you begin

You need the host's World Wide Port Names that you recorded in [Chapter 2](#) 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 `fcv 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 27.

For cfmode standby and mixed, special Fibre Channel switch zoning is required. See “[Zoning with cfmode standby and mixed](#)” on page 38.

- ◆ 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.
- ◆ 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 *Data ONTAP Block Access Management Guide* for your version of Data ONTAP.



## About this chapter

This chapter describes how to migrate an active/active storage system configuration to single\_image mode.

You can migrate to single\_image mode while installing the new FCP Linux Host Utilities, or as a stand-alone project. If you are migrating during the 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 28
- ◆ “[Completing prerequisites](#)” on page 30
- ◆ “[Upgrading to single\\_image mode](#)” on page 31

# Understanding cfmode

---

## What is single\_image mode?

In single\_image mode, both controllers in an active/active storage system configuration share the same World Wide Node Name (WWNN). Single\_image mode provides the following advantages:

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

Use the *Data ONTAP Block Access Management Guide for iSCSI and FCP* for your version of Data ONTAP to learn more about how single\_image mode works, but follow the procedures and recommendations in this document to migrate to single\_image mode.

## Resolving LUN mapping conflicts

Single\_image mode does not allow the following mapping conflicts:

- ◆ 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 might 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 ostype conflicts in the active/active configuration.

## 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 World Wide Port Names (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 (PTP) 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. You might want to change the cabling from the storage system to the fabric to optimize high availability and performance.

For detailed information about the cabling configurations supported and recommended examples for systems with onboard FCP ports, see the *FCP Configuration Guide* at

[http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/).

## Completing prerequisites

---

**Prerequisite tasks**      Complete the following tasks before you change your cfmode:

| Task                                      | Description                                                                                                                                                                                                                                                                                                                         | What to do                                                                                                                                                                                                     |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Save configuration information.           | <p>Enter commands that display information about the current configuration and save the output to the local disk.</p> <p>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.</p> | <p>Perform a complete backup of the system's data.</p> <hr/> <p><b>Note</b></p> <p>If you are using a native disk manager, then use the appropriate commands to obtain information about your disks.</p> <hr/> |
| Identify affected hosts and initiators.   | You must reconfigure and reboot any hosts and initiators that access LUNs on the systems you are upgrading.                                                                                                                                                                                                                         | <p>Enter the following command:</p> <pre>igroup show</pre>                                                                                                                                                     |
| Schedule downtime for the storage system. | Hosts cannot access LUNs while you change 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.                                                                                                                                      |

# Upgrading to single\_image mode

---

**Procedure Overview**

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 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 “ <a href="#">Migrating to single_image mode</a> ” on page 31.                                          |

**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. |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2    | <p>Login at the console of the systems you are upgrading, and enter the following command to identify and resolve LUN and igroup mapping conflicts.</p> <pre><b>lun config_check -S</b></pre> <p><b>Result:</b> Mappings that are displayed using the “*” character indicate an error, and also provide LUN and igroup information that you can use to resolve conflicts. Use the <code>lun map</code> command to re-map LUNS to a new ID that is unique across the storage system. You use the <code>igroup</code> command to change igroup otypes, if required.</p> |
| 3    | <p>Switch to advanced privileges by entering the following command:</p> <pre><b>priv set advanced</b></pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 4    | <p>Stop the FCP service on both systems. On each system, enter:</p> <pre><b>fcg stop</b></pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 5    | <p>Set the cfmode to single_image. On each system console, enter:</p> <pre><b>fcg set cfmode single_image</b></pre> <p>If you see messages indicating mapping conflicts, then return to <a href="#">Step 2</a> to resolve the mapping conflicts.</p>                                                                                                                                                                                                                                                                                                                  |
| 6    | <p>Start the FCP service on each system in the active/active configuration. On each system, enter the following command:</p> <pre><b>fcg start</b></pre>                                                                                                                                                                                                                                                                                                                                                                                                              |
| 7    | <p>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:</p> <pre><b>fcg nodename</b></pre>                                                                                                                                                                                                                                                                                                     |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2    | <p>Login at the console of the systems you are upgrading, and enter the following command to identify and resolve LUN and igroup mapping conflicts.</p> <pre><b>lun config_check -S</b></pre> <p><b>Result:</b> Mappings that are displayed using the “*” character indicate an error, and also provide LUN and igroup information that you can use to resolve conflicts. Use the <code>lun map</code> command to re-map LUNS to a new ID that is unique across the storage system. You use the <code>igroup</code> command to change igroup otypes, if required.</p> |
| 3    | <p>Switch to advanced privileges by entering the following command:</p> <pre><b>priv set advanced</b></pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 4    | <p>Stop the FCP service on both systems. On each system, enter:</p> <pre><b>fcg stop</b></pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 5    | <p>Set the cfmode to single_image. On each system console, enter:</p> <pre><b>fcg set cfmode single_image</b></pre> <p>If you see messages indicating mapping conflicts, then return to <a href="#">Step 2</a> to resolve the mapping conflicts.</p>                                                                                                                                                                                                                                                                                                                  |
| 6    | <p>Start the FCP service on each system in the active/active configuration. On each system, enter the following command:</p> <pre><b>fcg start</b></pre>                                                                                                                                                                                                                                                                                                                                                                                                              |
| 7    | <p>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:</p> <pre><b>fcg nodename</b></pre>                                                                                                                                                                                                                                                                                                     |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                         |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 8    | <p>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.</p> <pre><b>fcp config</b></pre> <p>For more information about zoning, see “<a href="#">Validating zoning and connectivity</a>” on page 34.</p> |
| 9    | <p>Return to administrative privileges. On each storage system, enter:</p> <pre><b>priv set admin</b></pre>                                                                                                                                                                                                                                                                    |

## Validating zoning and connectivity

**Changes to WWNNs and WWPNs:** After you change the cfmode setting to `single_image`, both controllers share a single World Wide Node Name. 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 portsets to reduce the number of paths in your configuration to comply with the number of paths supported by your multipathing software.

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

For details about portsets, see the *Data ONTAP 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 FCP Configuration Guide at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/).

**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.



**Before connecting the Linux host**

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.

**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.

# 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

See the FCP Configuration Guide at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/) for details of supported Fibre Channel configurations and switch zoning.

## 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 required HBA and driver parameters](#)” on page 19.

## 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 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 required HBA and driver parameters](#)” on page 19.

### Single\_image cfmode required

A direct Fibre Channel connection to an active/active storage system configuration requires that the storage system is using single\_image cfmode.

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

```
fcpx show cfmode
```

For information on changing to single\_image cfmode, see Chapter 4, “[Migrating a storage system to single\\_image mode](#),” on page 27.

### 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.



## Red Hat Enterprise Linux

This chapter applies only to Red Hat Enterprise Linux (RHEL) hosts. For SuSE Linux Enterprise Server (SLES) hosts, see Chapter 7, “[Accessing LUNs from a SLES host](#),” on page 67 instead.

## Starting FCP

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.

## Accessing multipath LUNs

If you are using the dm-multipath support, you access LUNs differently than if 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

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

---

## For detailed information

For detailed information, see the following sections:

- ◆ “[Loading the HBA driver on the host](#)” on page 42
- ◆ “[Editing the host’s /etc/multipath.conf file](#)” on page 43
- ◆ “[Configuring the multipath service to start automatically](#)” on page 47
- ◆ “[Accessing LUNs using dm-multipath](#)” on page 48
- ◆ “[Accessing LUNs without dm-multipath](#)” on page 56
- ◆ “[Stopping or reloading the HBA driver and shutting down](#)” on page 65

# Loading the HBA driver on the host

---

- 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 Fibre Channel protocol on a Linux host, load the driver for your HBAs.

If you are using dm-multipath, 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    | Entering the following command on the Linux host:<br><code>/sbin/modprobe -v qla2300</code> |

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

| Step | Action                                                                                   |
|------|------------------------------------------------------------------------------------------|
| 1    | Entering the following command on the Linux host:<br><code>/sbin/modprobe -v lpfc</code> |

# Editing the host's /etc/multipath.conf file

**Note**\_\_\_\_\_

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

\_\_\_\_\_

## Copying the file

The multipath.conf file contains configuration settings for multipathing. A sample multipath.conf file is copied to your host when you install the multipath-tools RPM. The file is:

`/usr/share/doc/packages/multipath-tools/multipath.conf.annotated`

Copy this file to `/etc/` and rename it `multipath.conf`.

## Editing the file

You need to edit the `/etc/multipath.conf` file to exclude (“blacklist”) the local hard drives and other resources that should not be included in the multipathing configuration.

To edit the `/etc/multipath.conf` file to blacklist local drives, complete the following steps:

| Step | Action                                                                                                                                                                                                                                                                                                                                           |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Open <code>/etc/multipath.conf</code> with a text editor.                                                                                                                                                                                                                                                                                        |
| 2    | Locate the following statement:<br><code>prio_callout "/sbin/mpath_prio_ontap/dev/%n"</code><br>Replace it with the following statement:<br><code>prio_callout "/opt/netapp/santools/mpath_prio_ontap/dev/%n"</code><br>If the statement is not found, add it to the <code>/etc/multipath.conf</code> file as shown in <a href="#">Step 10</a> . |
| 3    | Comment out the following line:<br><code>devnode_blacklist { devnode "*" }</code>                                                                                                                                                                                                                                                                |
| 4    | Delete the comment characters from the remaining <code>devnode_blacklist</code> command.                                                                                                                                                                                                                                                         |
| 5    | Locate the section labeled <code>## name : blacklist</code>                                                                                                                                                                                                                                                                                      |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6    | Delete the comment characters (#) to enable the code in the blacklist section.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 7    | If the command in the blacklist section is <code>blacklist</code> , rename the command <code>devnode_blacklist</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 8    | The default command excludes IDE hard drives <code>/dev/hd*</code> devices from <code>/dev/hda</code> through <code>/dev/hdz</code> . If your local drives are IDE drives in this range, no further changes to the blacklist section are needed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 9    | <p>If your local drives are SCSI drives, add a new line within the <code>devnode_blacklist</code> command to exclude the specific <code>/dev/sd*</code> devices. Be sure to blacklist only your local drives. Each path to an FCP LUN is also a <code>/dev/sd*</code> device that should be configured for multipathing.</p> <p>For example, if you have two local SCSI drives, <code>/dev/sda</code> and <code>/dev/sdb</code>, add the following line:</p> <pre>devnode "sd[a-b]\$" </pre> <p><b>Note</b></p> <p>Be sure to include the “\$” in the <code>devnode</code> statement when excluding local SCSI drives. Without the \$, the multipath code would exclude devices including <code>/dev/sdaa</code>, <code>/dev/sdab</code>, and <code>/dev/sdbb</code>, which will prevent multipathing from working if you have many LUN paths.</p> |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 10   | <p>If it does not already exist, create a device-specific section at the end of the file for the storage system as follows. These settings apply only to NetApp storage systems.</p> <pre> <b>devices {</b>   <b>device {</b>     <b>vendor</b>                <b>"NETAPP"</b>     <b>product</b>               <b>"LUN"</b>     <b>path_grouping_policy</b>  <b>group_by_prio</b>     <b>getuid_callout</b>       <b>"/sbin/scsi_id -g -u -s /block/%n"</b>     <b>prio_callout</b>         <b>"/opt/netapp/santools/mpath_</b> <b>prio_ontap /dev/%n"</b>     <b>features</b>              <b>"1 queue_if_no_path"</b>     <b>path_checker</b>         <b>readsector0</b>     <b>failback</b>             <b>immediate</b>   <b>}</b> <b>}</b> </pre> |
| 11   | Save the changes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

The following example shows the required changes to the `/etc/multipath.conf` file for a system with two local SCSI drives.

```
defaults {
    user_friendly_names yes
}
#
devnode_blacklist {
    devnode "sd[a-b]$"
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z]"
    devnode "^cciss!c[0-9]d[0-9]*"
}
#
devices {
    device {
        vendor                "NETAPP"
        product                "LUN"
        path_grouping_policy   group_by_prio
        getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
        prio_callout           "/opt/netapp/santools/mpath_
prio_ontap /dev/%n"
        features                "1 queue_if_no_path"
        path_checker            readsector0
        failback                immediate
    }
}
```

## Configuring the multipath service to start automatically

---

### Note

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

---

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

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

# Accessing LUNs using dm-multipath

---

## Understanding device persistence

The dm-multipath code creates persistent devices for each LUN in the `/dev/mapper/` directory on the Linux host. Instead of working with the `/dev/sdx` devices for multipath LUNs, you should use the devices in `/dev/mapper/`. You can create a file system directly on a multipath device in `/dev/mapper/`, or you can use the device directly as a raw device. There is no need to create a partition or label on the multipath device.

The devices in `/dev/mapper/` are also used for entries in the `/etc/fstab` file.

## Understanding automounting

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.

## For detailed information

For detailed information, see the following sections:

- ◆ [“Viewing a list of LUNs”](#) on page 49
- ◆ [“Viewing the dm-multipath configuration and multipath devices”](#) on page 51
- ◆ [“Creating and mounting a file system on a multipath device”](#) on page 52
- ◆ [“Accessing LUNs as a raw devices”](#) on page 54
- ◆ [“Tuning the dm-multipath configuration”](#) on page 55

## Viewing a list of LUNs

---

### Using sanlun

The `sanlun` command is installed in the `/opt/netapp/santools` directory as part of the FCP Linux Host Utilities. 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 mapped to the Linux host from each of two storage systems. There are two paths from the host to each storage system, which is why each LUN appears twice.

```
host:/ # sanlun lun show all
filer:   lun-pathname    device filename adapter protocol lun size      lun state
ss1:     /vol/vol1/linux_1 /dev/sde      host21  FCP       7g (7516192768)    GOOD
ss1:     /vol/vol1/linux_1 /dev/sdd      host20  FCP       7g (7516192768)    GOOD
ss2:     /vol/vol1/linux_2 /dev/sdc      host19  FCP       5g (5368709120)    GOOD
ss2:     /vol/vol1/linux_2 /dev/sdb      host18  FCP       5g (5368709120)    GOOD
```

Note that the `sanlun` command shows the mapping between the LUN path on the storage system and the `/dev/sdx` device on the Linux host. In this example, storage system `ss1` LUN `/vol/vol1/linux_1` is seen as both `/dev/sde` and `/dev/sdd` on the host.

### Discovering new LUNs

If you create a new LUN and map it to the Linux host, you can discover the LUN by reloading the HBA drivers on the host.

To rescan the HBAs for new LUNs without reloading the drivers, complete one of the following steps:

**Rescanning QLogic HBAs:** To rescan QLogic HBAs to discover new LUNs and targets, run the `qla2xxx_lun_rescan` script included with the Host Utilities on the Linux host:

```
/opt/netapp/santools/qla2xxx_lun_rescan all
```

**Rescanning Emulex HBAs:** To rescan Emulex HBAs to discover new LUNs and targets, run the Emulex `lun_scan` script on the Linux host:

```
/usr/sbin/lpfc/lun_scan all
```

You can use the `sanlun` command to verify that the new LUNs were discovered.

Once discovered, the LUNs should be automatically added to the dm-multipath configuration. Run the `multipath -v2` command to update the multipath configuration manually if needed.

## Viewing the dm-multipath configuration and multipath devices

---

### Viewing 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.

### Viewing the multipath devices

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

```
multipath -d -l
```

**Example:** The following output shows two LUNs, with two paths each.

```
mpath1 (360a9800043346536534a344a794e6d59)
[size=7 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [active]
  \_ 20:0:0:0 sdd 8:48 [active][ready]
  \_ 21:0:0:0 sde 8:64 [active][ready]

mpath0 (360a9800043346536684a34425579312d)
[size=5 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [active]
  \_ 18:0:0:0 sdb 8:16 [active][ready]
  \_ 19:0:0:0 sdc 8:32 [active][ready]
```

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

To view the `/dev/mapper` devices, enter **`ls -l /dev/mapper/`**.

**Example:** The following example shows two devices, `mpath0` and `mpath1`.

```
brw-rw---- 1 root disk 253,  3 Apr 14 11:44 mpath0
brw-rw---- 1 root disk 253,  2 Apr 14 11:44 mpath1
```

## 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. If you are using the LUN as a raw device, skip this section.

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, complete the following step:

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/mpath0
```

### Mounting the file system

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

To mount the file system, complete the following steps:

| Step | Action                                                               |
|------|----------------------------------------------------------------------|
| 1    | Create a mount point directory using the <code>mkdir</code> command. |
| 2    | Open the <code>/etc/fstab</code> file with an editor.                |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | <p>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):</p> <pre><i>device mount_point type defaults 0 0</i></pre> <p><i>device</i> is the name of the device in the /dev/mpath/ directory</p> <p><i>mount_point</i> is the mount point you created for the file system</p> <p><i>type</i> is the file system type, such as ext2 or ext3</p> <p><b>Example:</b></p> <pre>/dev/mapper/mpath0 /mnt/ss1/lun_0 ext3 defaults 0 0</pre> |
| 4    | <p>Enter <b>mount <i>mountpoint</i></b> to mount the file system.</p> <p><b>Example:</b></p> <pre>mount /mnt/ss1/lun_0</pre>                                                                                                                                                                                                                                                                                                                                                                                        |
| 5    | <p>Enter <b>df</b> to display the mounted file systems. Verify the new file system is listed.</p> <p>Another option is to change to the mount point directory and try writing and reading a file to verify the file system mounted successfully.</p>                                                                                                                                                                                                                                                                |

## **Accessing LUNs as a raw devices**

---

### **Overview**

Instead of creating a file system, you can access a LUN as a raw device.

Be sure to use the `/dev/mapper/` device created by dm-multipath to access the raw device.

For the Linux 2.6 kernel used by RHEL 4, do not use the raw interface. Instead open the device using the open call with the `O_DIRECT` flag.

## 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 per 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 as it really depends on the workload and other factors such as the host memory and CPU speed.

The rr\_min\_io value can only be set by changing the defaults section in /etc/multipath.conf. You cannot change it in the devices section. You must put a section such as the following at the very top of the /etc/multipath.conf file:

```
defaults {
    rr_min_io      128
}
```

# Accessing LUNs without dm-multipath

---

## Limited configuration without dm-multipath

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 have 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 38. 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 required HBA and driver parameters”](#) on page 19.

Configurations without dm-multipath may have additional restrictions on host OS versions and HBA models. See the *NetApp FCP Host Compatibility Matrix* at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/).

## Understanding device persistence

LUNs on the storage system are seen by the Linux host as SCSI devices. Linux assigns an identifier, such as /dev/sda or /dev/sdb, to each LUN. The identifiers 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 may be mapped to a different identifier after a Linux reboot or a restart 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.

## 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 one of the following two tasks:

- ◆ If you are creating a file system on the LUN:

- ❖ 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.
- ❖ Create a file system on the partition.
- ❖ Create a label for the file system and add an entry for the file system's label in `/etc/fstab`.
- ◆ Configure the LUN as a raw device.

## For detailed information

For detailed information, see the following sections:

- ◆ [“Creating a partition”](#) on page 58
- ◆ [“Creating a file system”](#) on page 60
- ◆ [“Labeling and mounting the file system”](#) on page 61
- ◆ [“Accessing LUNs as raw devices”](#) on page 63
- ◆ [“Viewing LUN information”](#) on page 64

# 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, you still need to create a partition.

**Before you begin**      You should have one or more LUNs available to the host. Log into 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:

| Step | Action                                                                                                     |
|------|------------------------------------------------------------------------------------------------------------|
| 1    | Enter <b>/sbin/fdisk device</b><br><i>device</i> is the device name assigned to the LUN, such as /dev/sda. |
| 2    | Enter the command <b>n</b> to create a new partition.                                                      |
| 3    | Enter <b>p</b> to create a primary partition                                                               |
| 4    | Enter a partition number. For the first partition, enter <b>1</b> .                                        |
| 5    | Press Enter twice to accept the default cylinder values.                                                   |

| Step      | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |        |       |          |     |        |    |        |           |  |   |       |          |    |       |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------|-------|----------|-----|--------|----|--------|-----------|--|---|-------|----------|----|-------|
| 6         | <p>Enter <b>p</b> to print the partition information.</p> <p><b>Result:</b></p> <p>Command (m for help): p</p> <p>Disk /dev/sdb: 64 heads, <b>32 sectors</b>, 10240 cylinders<br/>Units = cylinders of 2048 * 512 bytes</p> <table><thead><tr><th>Device</th><th>Boot</th><th>Start</th><th>End</th><th>Blocks</th><th>Id</th><th>System</th></tr></thead><tbody><tr><td>/dev/sdb1</td><td></td><td>1</td><td>10240</td><td>10485744</td><td>83</td><td>Linux</td></tr></tbody></table> | Device | Boot  | Start    | End | Blocks | Id | System | /dev/sdb1 |  | 1 | 10240 | 10485744 | 83 | Linux |
| Device    | Boot                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Start  | End   | Blocks   | Id  | System |    |        |           |  |   |       |          |    |       |
| /dev/sdb1 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 1      | 10240 | 10485744 | 83  | Linux  |    |        |           |  |   |       |          |    |       |
| 7         | <p>If the sectors value from the previous step is not a multiple of 8, you must set the sectors manually to ensure best performance. Enter <b>q</b> to quit without creating the new partition. Refer to bug number 156121 on Bugs Online at <a href="http://now.netapp.com/NOW/cgi-bin/bol">http://now.netapp.com/NOW/cgi-bin/bol</a> for the latest information and specific fdisk parameters.</p>                                                                                    |        |       |          |     |        |    |        |           |  |   |       |          |    |       |
| 8         | <p>If the sector value is a multiple of 8, enter <b>w</b> to save the changes and exit fdisk.</p>                                                                                                                                                                                                                                                                                                                                                                                       |        |       |          |     |        |    |        |           |  |   |       |          |    |       |

# 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. These instructions show how to create a journaling ext3 file system, but you can create any type you like.

If your application needs a raw device, do not create a file system. Instead, see [“Accessing LUNs as raw devices”](#) on page 63.

**Using mkfs to create a file system**      To create a file system on the partition you created, run the `mkfs` command as follows:‘

| Step | Action                                                                                                                                                                   |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Enter <code>mkfs -t ext3 device</code><br><br>Where <i>device</i> is the partition. For example <code>/dev/sda1</code> is the first partition on <code>/dev/sda</code> . |

Accessing LUNs without dm-multipath

# Labeling and mounting the file system

---

About labels

A file system label creates a persistent mapping between a LUN and a file system. See “[Understanding device persistence](#)” on page 56 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    | <p>Enter <code>/sbin/e2label device label</code></p> <p><i>device</i> is the partition where you created the file system.</p> <p><i>label</i> is a string (maximum 16 characters) used to identify the file system.</p> <p><b>Example:</b> <code>/sbin/e2label /dev/sda1 lun0_label</code></p> |

Mounting the file system

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

| Step | Action                                                                                                      |
|------|-------------------------------------------------------------------------------------------------------------|
| 1    | <p>Create a mount point for the file system.</p> <p><b>Example:</b> <code>mkdir /mnt/filer0_lun0</code></p> |
| 2    | <p>Open the file /etc/fstab using a text editor.</p>                                                        |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | <p>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):</p> <pre>LABEL=<i>label</i> <i>mount_point</i> auto defaults 0 0</pre> <p><i>label</i> is the file system label you created above</p> <p><i>mount_point</i> is the mount point you created for the file system</p> <p><b>Example:</b></p> <pre>LABEL=lun0_label /mnt/filer0_lun0 auto defaults 0 0</pre> |
| 4    | Save your changes to /etc/fstab and exit the editor.                                                                                                                                                                                                                                                                                                                                                                                              |
| 5    | Enter <b>mount LABEL=<i>label</i></b> to mount the file system. The file system will be automatically mounted whenever the Linux host is rebooted.                                                                                                                                                                                                                                                                                                |
| 6    | <p>Enter <b>df</b> to display the mounted file systems. Verify the new file system is listed.</p> <p>Another option is to change to the mount point directory and try writing and reading a file to verify the file system mounted successfully.</p>                                                                                                                                                                                              |

## **Accessing LUNs as raw devices**

---

Instead of creating a file system, you can access a LUN as a raw device.

For the Linux 2.6 kernel used by RHEL 4, do not use the raw interface. Instead open the device using the open call with the O\_DIRECT flag.

Accessing LUNs without dm-multipath

## Viewing LUN information

---

**Viewing information about LUNs**

You can view information about LUNs in the following ways:

- ◆ 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 the FCP Linux Host Utilities.

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

# Stopping or reloading the HBA driver and shutting down

---

## Stopping the HBA driver

Before you stop or reload the HBA driver, unmount any file systems on Fibre Channel LUNs, flush the multipath maps, and stop the dm-multipath daemon.

To flush the multipath maps, enter the following command:

```
multipath -F
```

To stop the multipath daemon, enter the following command:

```
/etc/init.d/multipathd stop
```

Unload all of the dm-multipath modules. Depending on which modules are loaded, the command should be similar to:

```
modprobe -r dm_round_robin dm_mirror dm_multipath dm_mod
```

After stopping multipathing, unload the HBA driver.

**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
```

## Reloading the HBA driver

To reload the HBA driver, first complete the steps above for stopping the driver. Then enter the following commands to restart the HBA driver and dm-multipath. Be sure to reload specific dm-multipath modules that were previously loaded.

```
modprobe qla2300 (for QLogic 2300 series HBA)
```

```
modprobe qla2400 (for QLogic 2400 series HBA)
```

```
modprobe lpfc (for Emulex HBA)
```

```
modprobe dm_mod
```

```
modprobe dm_multipath
```

```
modprobe dm_mirror
```

```
modprobe dm_round_robin
```

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

```
multipath -v 2
```

## 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
```

## SuSE Linux Enterprise Server

This chapter applies only to SuSE Linux Enterprise Server (SLES) hosts. For Red Hat Enterprise Linux (RHEL) hosts, see Chapter 6, “[Accessing LUNs from a RHEL host](#),” on page 41 instead.

## Loading the HBA 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.

## Overview of accessing LUNs

The steps for accessing LUNs depend 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 68
- ◆ “[Editing the host’s /etc/multipath.conf file](#)” on page 69
- ◆ “[Starting the multipath service](#)” on page 73
- ◆ “[Configuring the multipath service to start automatically](#)” on page 74
- ◆ “[Accessing LUNs with dm-multipath support](#)” on page 75
- ◆ “[Accessing LUNs without dm-multipath](#)” on page 84
- ◆ “[Stopping or reloading the HBA driver and shutting down](#)” on page 92

# Loading the HBA driver on the host

---

- 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 Fibre Channel protocol on a Linux host, load the driver for your HBAs.

If you are using dm-multipath, 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    | Entering the following command on the Linux host:<br><br><code>/sbin/modprobe -v qla2300</code> |

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

| Step | Action                                                                                       |
|------|----------------------------------------------------------------------------------------------|
| 1    | Entering the following command on the Linux host:<br><br><code>/sbin/modprobe -v lpfc</code> |

# Editing the host's /etc/multipath.conf file

**Note**\_\_\_\_\_

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

\_\_\_\_\_

## Copying the file

The multipath.conf file contains configuration settings for multipathing. A sample multipath.conf file is copied to your host when you install the multipath-tools RPM. The file is:

`/usr/share/doc/packages/multipath-tools/multipath.conf.annotated`

Copy this file to `/etc/` and rename it `multipath.conf`.

## Editing the file

You need to edit the `/etc/multipath.conf` file to exclude (“blacklist”) the local hard drives and other resources that should not be included in the multipathing configuration.

To edit the `/etc/multipath.conf` file to blacklist local drives, complete the following steps:

| Step | Action                                                                                                                                                                                                                                                                                                                                          |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Open <code>/etc/multipath.conf</code> with a text editor.                                                                                                                                                                                                                                                                                       |
| 2    | Locate the following statement:<br><code>prio_callout "/sbin/mpath_prio_ontap/dev/%n"</code><br>Replace it with the following statement:<br><code>prio_callout "/opt/netapp/santools/mpath_prio_ontap/dev/%n"</code><br>If the statement is not found, add it to the <code>/etc/multipath.conf</code> file as shown in <a href="#">Step 8</a> . |
| 3    | Locate the section labeled <code>## name : blacklist</code>                                                                                                                                                                                                                                                                                     |
| 4    | Delete the comment characters ( <code>#</code> ) to enable the code in the blacklist section.                                                                                                                                                                                                                                                   |
| 5    | If the command in the blacklist section is <code>blacklist</code> , rename the command <code>devnode_blacklist</code> .                                                                                                                                                                                                                         |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6    | <p>The default command excludes IDE hard drives /dev/hd<math>x</math> devices from /dev/hda through /dev/hdz. However, there may be a syntax error in the sample <code>/etc/multipath.conf</code> file that prevents local drives from being properly blacklisted.</p> <p>If the following line has two extra bracket characters (bold):</p> <pre>devnode "^hd[a-z] <b>[</b>[0-9]*<b>]</b>"</pre> <p>Delete the bracket so the line reads as follows:</p> <pre>devnode "^hd[a-z] [0-9]*"</pre> <p>If the following line has two extra bracket characters (bold):</p> <pre>devnode "^cciss!c[0-9]d[0-9]*<b>[</b>p[0-9]*<b>]</b>"</pre> <p>Delete the brackets so the line reads as follows:</p> <pre>devnode "^cciss!c[0-9]d[0-9]*p[0-9]*"</pre> <p>If your local drives are IDE drives in the above range, no further changes are needed; save the file and quit.</p> |
| 7    | <p>If your local drives are SCSI drives, add a new line within the <code>devnode_blacklist</code> command to exclude the specific /dev/sd<math>x</math> devices. Be sure to blacklist only your local drives. Each path to an FCP LUN is also a /dev/sd<math>x</math> device that should be configured for multipathing.</p> <p>For example, if you have two local SCSI drives, /dev/sda and /dev/sdb, add the following line:</p> <pre>devnode "sd[a-b]\$"</pre> <p><b>Note</b></p> <p>Be sure to include the "\$" in the <code>devnode</code> statement when excluding local SCSI drives. Without the \$, the multipath code would exclude devices including /dev/sdaa, /dev/sdab, and /dev/sdbb, which will prevent multipathing from working if you have many LUN paths.</p>                                                                                      |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 8    | <p>Create a device-specific section at the end of the file for the storage system as follows. These settings apply only to NetApp storage systems.</p> <pre> <b>devices {</b>   <b>device {</b>     <b>vendor</b>                "NETAPP"     <b>product</b>               "LUN"     <b>path_grouping_policy</b>  <b>group_by_prio</b>     <b>getuid_callout</b>       "/sbin/scsi_id -g -u -s /block/%n"     <b>prio_callout</b>         "/opt/netapp/santools/mpath_ <b>prio_ontap</b> /dev/%n"     <b>features</b>              "1 queue_if_no_path"     <b>path_checker</b>         readsector0     <b>failback</b>              <b>immediate</b>   } <b>}</b> </pre> |
| 9    | Save the changes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

The following example shows the required changes to the `/etc/multipath.conf` file for a system with two local SCSI drives.

```

##
## name      : blacklist
## scope     : multipath & multipathd
## desc      : list of device names to discard as not multipath
candidates
## default   : cciss, fd, hd, md, dm, sr, scd, st, ram, raw, loop
##
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]*"
    devnode "sd[a-b]$"
}
#
##
devices {
    device {
        vendor                "NETAPP"
        product               "LUN"
        path_grouping_policy   group_by_prio
        getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"

```

```
        prio_callout "/opt/netapp/santools/mpath_  
prio_ontap /dev/%n"  
        features "1 queue_if_no_path"  
        path_checker readsector0  
        failback immediate  
    }  
}
```

# Starting the multipath service

---

**Note**\_\_\_\_\_

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

---

Once you have started the HBA drivers, start the multipath service on the Linux host. To start the multipath service, enter the following commands on the Linux host console:

| Step | Action                                                                                                                                                           |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Start the scan for multipath LUNs by entering the following command:<br><br><code>/etc/init.d/boot.multipath start</code>                                        |
| 2    | After the command in the previous step completes, start the multipath daemon by entering the following command:<br><br><code>/etc/init.d/multipathd start</code> |

# Configuring the multipath service to start automatically

---

**Note**\_\_\_\_\_

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

---

Once you have started the multipath service on the Linux host, you can configure it to start automatically after reboot. To integrate the multipath setup into the boot sequence, enter the following command on the Linux host console:

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

# Accessing LUNs with dm-multipath support

---

## Overview

The process for accessing a LUN using dm-multipath is as follows:

- ◆ “[Verifying the dm-multipath policy and mapping](#)” on page 76
- ◆ “[Viewing a list of LUNs](#)” on page 77
- ◆ “[Creating and mounting a file system on a LUN](#)” on page 79
- ◆ “[Accessing LUNs as a raw devices](#)” on page 81
- ◆ “[Tuning the dm-multipath configuration](#)” on page 82
- ◆ “[Stopping dm-multipath](#)” on page 83

## **Verifying the dm-multipath policy and mapping**

---

### **Overview**

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

### **Verifying the configuration**

Enter the following command to display all of the information about the current dm-multipath configuration:

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

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

### **Updating the configuration**

To update the dm-multipath maps, enter the following command:

```
multipath -v2
```

## Viewing a list of LUNs

---

### Using sanlun

The `sanlun` command is installed by the Host Utilities in the `/opt/netapp/santools` directory. 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 mapped to the Linux host from each of two storage systems. There are two paths from the host to each storage system, which is why each LUN appears twice.

```
host:/ # sanlun lun show all
filer:  lun-pathname      device filename  adapter  protocol  lun size    lun state
ss1:    /vol/vol1/linux_1  /dev/sdc       host0    FCP       5g (5368709120)  GOOD
ss2:    /vol/vol1/linux_2  /dev/sdd       host0    FCP       6g (6442450944)  GOOD
ss1:    /vol/vol1/linux_1  /dev/sde       host1    FCP       5g (5368709120)  GOOD
ss2:    /vol/vol1/linux_2  /dev/sdf       host1    FCP       6g (6442450944)  GOOD
```

Note that the `sanlun` command shows the mapping between the LUN path on the storage system and the `/dev/sdx` device on the Linux host. In this example, storage system `ss1` LUN `/vol/vol1/linux_1` is seen as both `/dev/sdc` and `/dev/sde` on the host.

### Discovering new LUNs

If you create a new LUN and map it to the Linux host, you can discover the LUN by rescanning the HBA's targets.

**Rescanning QLogic HBAs:** To rescan QLogic HBAs to discover new LUNs and targets, run the `qla2xxx_lun_rescan` script included with the Host Utilities on the Linux host:

```
/opt/netapp/santools/qla2xxx_lun_rescan all
```

**Rescanning Emulex HBAs:** To rescan Emulex HBAs to discover new LUNs and targets, run the `Emulex lun_scan` script on the Linux host:

```
/usr/sbin/lpfc/lun_scan all
```

Use the `sanlun` command to verify that the new LUNs were discovered.

## Adding new LUNs to the dm-multipath configuration

Once you have discovered new LUNs, you need to add them to the dm-multipath configuration. To add LUNs to the multipath configuration, enter the following command on the Linux host console:

**multipath**

The output of the multipath command displays the configuration of the new LUN. The following example output shows one new LUN with two paths:

```
create: 360a9800043346536684a343369697562
[size=4 GB][features="0"][hwhandler="0"]
\_ round-robin 0 [prio=2]
  \_ 5:0:0:1      sdg 8:96  [ready]
  \_ 5:1:0:1      sdh 8:112 [ready]
```

Accessing LUNs with dm-multipath support

# Creating and mounting a file system on a LUN

---

**Overview**

Once a LUN is added to the multipath configuration, you can create a file system. If you are using the LUN as a raw device, skip this section.

When creating a file system, you use the multipath device in the `/dev/disk/by-name/` directory.

**Creating the file system**

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

Enter the following command on the Linux host console:

**mkfs -t *type* /dev/disk/by-name/*device***

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

*device* is the multipath device name of the LUN in the `/dev/disk/by-name/` directory.

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

```
mkfs -t ext3 /dev/disk/by-name/360a9800043346536684a337a7553424f
```

**Mounting the file system**

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

To mount the file system, complete the following steps:

| Step | Action                                                               |
|------|----------------------------------------------------------------------|
| 1    | Create a mount point directory using the <code>mkdir</code> command. |
| 2    | Open the <code>/etc/fstab</code> file with an editor.                |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | <p>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):</p> <pre>device mount_point type defaults 0 0</pre> <p><i>device</i> is the name in the /dev/disk/by-name/ directory</p> <p><i>mount_point</i> is the mount point you created for the file system</p> <p><i>type</i> is the file system type, such as ext2, ext3, or reiserfs</p> <p><b>Example:</b></p> <pre>/dev/disk/by-name/360a9800043346536684a337a7553424f /mnt/ss1/lun_0      ext3      defaults          0 0</pre> |
| 4    | Enter <b>mount -a</b> to mount the file system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 5    | <p>Enter <b>df</b> to display the mounted file systems. Verify the new file system is listed.</p> <p>Another option is to change to the mount point directory and try writing and reading a file to verify the file system mounted successfully.</p>                                                                                                                                                                                                                                                                                                                 |

## **Accessing LUNs as a raw devices**

---

### **Overview**

Instead of creating a file system, you can access a LUN as a raw device.

In 2.6 kernels and later, the `raw` command is not recommended. For SLES 9, if you are writing scripts that access the LUNs as raw devices, you open the SCSI disk devices using the `open` call with the `O_DIRECT` flag.

Be sure to use the `/dev/disk/by-name/name` path created by `dm-multipath` to access the raw device.

## 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.

**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 as it really depends on the workload and other factors such as the host memory and CPU speed.

The `rr_min_io` value can only be set by changing the `defaults` section in `/etc/multipath.conf`. You cannot change it in the `devices` section. You must put a section such as the following at the very top of the `/etc/multipath.conf` file:

```
defaults {  
    rr_min_io      128  
}
```

# Stopping dm-multipath

---

Overview

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.

Stopping the services

To stop dm-multipath and FCP on the host, complete the following steps:

| Step | Action                                                                                                                                                                                  |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Unmount all file systems on FCP LUNs.                                                                                                                                                   |
| 2    | Flush the multipath maps by entering the following command:<br><code>multipath -F</code>                                                                                                |
| 3    | Stop the multipath daemon by entering the following command:<br><code>/etc/init.d/multipathd stop</code>                                                                                |
| 4    | Unload all of the dm-multipath modules by entering the following command:<br><code>/etc/init.d/boot.multipath stop</code>                                                               |
| 5    | Stop the HBA driver by entering the following command.<br><br>For QLogic HBAs:<br><code>/sbin/modprobe -r qla2300</code><br><br>For Emulex HBAs:<br><code>/sbin/modprobe -r lpfc</code> |

# Accessing LUNs without dm-multipath

---

## Limited configuration without dm-multipath

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 have 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 38. 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 required HBA and driver parameters”](#) on page 19.

Configurations without dm-multipath may have additional restrictions on host OS versions and HBA models. See the *NetApp FCP Host Compatibility Matrix* at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/).

## How LUNs are integrated into the Linux file system

Linux uses symbolic links to provide host applications with reliable mappings between device nodes and LUNs. Linux assigns SCSI device nodes dynamically when LUNs are detected. This means that the mapping of SCSI devices to LUNs might be inconsistent each time the HBA driver is loaded, and SCSI commands might be sent to the wrong storage system or LUN. If you have applications that access storage system LUNs, you must ensure that these applications use the symbolic links instead of the standard SCSI device nodes.

The FCP LUN devices are under the `/dev/disk/by-id/` directory.

The following shows example shows three FCP devices in the `/dev/disk/by-id/` directory.

```
lrwxrwxrwx 1 root root 9 May 26 11:07 360a98000433464624f4a34732f785a47 -> ../../sda
lrwxrwxrwx 1 root root 9 May 26 11:07 360a98000433464624f4a353036375076 -> ../../sdb
lrwxrwxrwx 1 root root 9 May 26 11:07 360a98000433464624f4a3531677a3071 -> ../../sdc
```

## Viewing information about LUNs

You can view information about LUNs in the following ways:

- ◆ **Symbolic links**—The `ls -l` command lists the symbolic links the initiator creates under the `/dev/disk/by-id` directory.
- ◆ **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.

**Viewing symbolic links:** To view symbolic links and the storage system node name to which they are mapped, complete the following step:

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>On the SLES host console, enter the following command:</p> <pre>ls -l /dev/disk/by-id</pre> <p><b>Example:</b></p> <pre># ls -l /dev/disk/by-id drwxr-xr-x  2 root root   4096 Jul 13 15:00 . drwxr-xr-x  5 root root   4096 Jun 11 21:38 .. lrwxrwxrwx  1 root root          9 May 26 11:07 360a98000433464624f4a34732f785a47 -&gt; ../../../../sda lrwxrwxrwx  1 root root          9 May 26 11:07 360a98000433464624f4a353036375076 -&gt; ../../../../sdb lrwxrwxrwx  1 root root          9 May 26 11:07 360a98000433464624f4a3531677a3071 -&gt; ../../../../sdc</pre> |

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

| Step | Action                                                                  |
|------|-------------------------------------------------------------------------|
| 1    | On the host, change to the <code>/opt/netapp/santools</code> directory. |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2    | <p>Enter the following command:</p> <pre><b>sanlun lun show all</b></pre> <p><b>Result:</b> The <code>sanlun</code> command displays information about FCP devices and the storage system LUNs they are mapped to, including LUN size and state.</p> <p><b>Example:</b> The following lines show sample <code>sanlun</code> command output.</p> <pre>[root@host_1 /]# /opt/netapp/santools/sanlun lun show all filer:  lun-pathname      device filename adapter protocol lun size      lun state ssl:   /vol/luns/lun0     /dev/sda          host0    FCP      10g (10737418240)  GOOD ssl:   /vol/luns/lun1     /dev/sdb          host0    FCP      10g (10737418240)  GOOD ssl:   /vol/luns/lun2     /dev/sdc          host1    FCP      10g (10737418240)  GOOD</pre> |

## Accessing LUNs

You can access LUNs on the storage system in one of the following ways:

- ◆ Creating a file system on the LUN—For details, see “[Creating a file system on a LUN](#)” on page 88.
- ◆ Accessing the LUN as a raw device—For details, see “[Accessing LUNs as raw devices](#)” on page 91.

# Creating a file system on a LUN

---

If you are not using raw access, you must create a partition and file system on the LUN and automount that file system.

## Using fdisk to create a partition

To create a partition of a LUN:

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

## Creating a file system

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>On the host console, enter the following command:</p> <pre><b>mkfs -V -t <i>filesystem_type</i> <i>partition</i></b></pre> <p><i>partition</i> is the name of the partition you just created.</p> <p>For SLES 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.</p> <p><i>filesystem_type</i> is the type of file system you want to create.</p> <p><b>Example:</b> The following example creates an ext3 file system on the first partition of LUN 0.</p> <pre>mkfs -V -t ext3 /dev/disk/by-id/360a98000433464624f4a34732f785a47-0p1</pre> |

## Mount the file system

To mount the file system, complete the following steps:

| Step | Action                                                                                                                                                                                                                                                                                                                             |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | On the host, create a mount point for the LUNs.                                                                                                                                                                                                                                                                                    |
| 2    | With a text editor, open the host's <code>/etc/fstab</code> file.                                                                                                                                                                                                                                                                  |
| 3    | <p>Add an entry for the file system you created.</p> <p><b>Example:</b></p> <pre>/dev/disk/by-id/360a98000433464624f4a34732f785a47-0p1 /mnt_pt ext3 defaults 0 0</pre> <p><i>mnt_pt</i> is the mount point you created in <a href="#">Step 1</a>.</p> <p><b>Result:</b> The file systems listed are mounted when Linux starts.</p> |
| 4    | Save your changes and exit the <code>/etc/fstab</code> file.                                                                                                                                                                                                                                                                       |
| 5    | <p>On the host console, enter the following command to mount the file systems you added:</p> <pre>mount</pre>                                                                                                                                                                                                                      |

## **Accessing LUNs as raw devices**

---

### **SLES 9**

In 2.6 kernels and later, the `raw` command is not recommended. For SLES 9, if you are writing scripts that access the LUNs as raw devices, you open the SCSI disk devices using the `open` call with the `O_DIRECT` flag.

# Stopping or reloading the HBA driver and shutting down

---

## Stopping the HBA driver

Before you stop or reload the HBA driver, unmount any file systems on Fibre Channel LUNs, flush the multipath maps, and stop the dm-multipath daemon.

To flush the multipath maps, enter the following command:

```
multipath -F
```

To stop the multipath daemon, enter the following command:

```
/etc/init.d/multipathd stop
```

Unload all of the dm-multipath modules by entering the following command:

```
/etc/init.d/boot.multipath stop
```

After stopping multipathing, unload the HBA driver.

**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
```

## Reloading the HBA driver

To reload the HBA driver, first complete the steps above for stopping the driver. Then enter the following commands to restart the HBA driver and dm-multipath. Be sure to reload specific dm-multipath modules that were previously loaded.

```
modprobe qla2300 (for QLogic 2300 series HBA)
```

```
modprobe qla2400 (for QLogic 2400 series HBA)
```

```
modprobe lpfc (for Emulex HBA)
```

```
/etc/init.d/boot.multipath start
```

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

```
multipath -v 2
```

## Shutting down the Linux host

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

```
multipath -F
```

## Note

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

## 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. To rescan the HBAs, run the rescan script for your HBA model as follows:

**QLogic:** `/opt/netapp/santools/qla2xxx_lun_rescan all`

**Emulex:** `/usr/sbin/lpfc/lun_scan all`

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:<br><br><a href="http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/">http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/</a> |
| Cabling               | Verify that the cables between the host and the storage system are properly connected.                                                                                                                                                                                                                                                      |
| FCP service status    | Verify that the FCP service is licensed and started on the storage system according to the procedure described in the <i>Data ONTAP Block Access Management Guide</i> .                                                                                                                                                                     |

| Configuration setting     | What to do                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Initiator login           | <p>Verify that the initiator is logged in to the storage system by entering the <code>fcpl show initiator</code> command on the storage system console.</p> <p>If the initiator is configured and logged in to the storage system, the storage system console displays the initiator WWPN.</p>                                                                                                                                                                                                          |
| HBA World Wide Port Names | <p>Verify that you are using the correct initiator WWPNs in the igroup configuration.</p> <p>On the storage system, use the <code>igroup show</code> 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.</p>                                                                                                                         |
| LUN mappings              | <p>Verify that the LUNs are mapped to an igroup.</p> <p>On the storage system, use one of the following commands:</p> <p><code>lun show -m</code> displays all LUNs and the igroups to which they are mapped.</p> <p><code>lun show -g <i>igroup-name</i></code> displays the LUNs mapped to a specific igroup.</p>                                                                                                                                                                                     |
| System requirements       | <p>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 FCP Host Compatibility Matrix at the following URL:</p> <p><a href="http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/">http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/</a></p> |

## 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 the `devnode_blacklist` section of the `/etc/multipath.conf` file. Pay special attention to any statements added to blacklist local SCSI disks. Note that `devnode "sd[a-b]"` excludes not only `/dev/sda` and `/dev/sdb`, but also `/dev/sdaa`, `/dev/sdba`, and so on. Be sure to use `devnode "sd[a-b]$"` instead.

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 201302 on Bugs Online at <http://now.netapp.com/NOW/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 may 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. Refer to bug number 156121 on Bugs Online at <http://now.netapp.com/NOW/cgi-bin/bol> for specific steps.

## 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 HBA driver and applications](#)” on page 17.

## **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. See “[Setting required HBA and driver parameters](#)” on page 19.

Verify that you have correctly zoned your Fibre Channel switches. See “[Switched fabric configurations](#)” on page 38.

For multipath LUN, 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
```