

# FCP Linux® Host Utilities 4.0 for Native OS and Veritas®

## **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 215-03003\_A0  
November 2007

# Copyright and trademark information

---

## Copyright information

Copyright © 1994–2007 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, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, SyncMirror, Topio, VFM, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. Cryptainer, Cryptoshred, Datafort, and Decru are registered trademarks, and Lifetime Key Management and OpenKey are trademarks, of Decru, a Network Appliance, Inc. company, 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. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexShare, FPolicy, HyperSAN, InfoFabric, LockVault, NOW, NOW NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simplicore, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, SohoFiler, SpinMirror, SpinRestore, SpinShot, SpinStor,

StoreVault, vFiler, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the U.S.A.

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

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

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

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> . . . . .                                              | ix |
| <b>Chapter 1</b> | <b>Overview</b> . . . . .                                             | 1  |
|                  | The setup procedure . . . . .                                         | 2  |
|                  | Where to find the information you need . . . . .                      | 5  |
|                  | Understanding the Host Utilities components and requirements. . . . . | 9  |
|                  | Understanding host and storage system communication . . . . .         | 12 |
| <b>Chapter 2</b> | <b>Installing HBAs and Software on the Host</b> . . . . .             | 15 |
|                  | Installing the FCP Linux Host Utilities software. . . . .             | 16 |
|                  | Removing the previous HBA drivers. . . . .                            | 18 |
|                  | Verifying or installing the required multipathing packages . . . . .  | 20 |
|                  | Installing HBAs. . . . .                                              | 21 |
|                  | Downloading and installing the HBA driver and applications . . . . .  | 22 |
|                  | Setting required HBA and driver parameters. . . . .                   | 24 |
|                  | Recording the WWPNs of the HBAs. . . . .                              | 26 |
| <b>Chapter 3</b> | <b>Configuring the Storage System for FCP</b> . . . . .               | 29 |
| <b>Chapter 4</b> | <b>Migrating a storage system to single_image mode</b> . . . . .      | 31 |
|                  | Understanding cfmode . . . . .                                        | 32 |
|                  | Migrating to single_image mode . . . . .                              | 34 |
| <b>Chapter 5</b> | <b>Configuring FCP cabling and switches</b> . . . . .                 | 39 |
|                  | Switched fabric configurations . . . . .                              | 40 |
|                  | Direct-connected configurations . . . . .                             | 42 |

|                      |                                                                              |               |
|----------------------|------------------------------------------------------------------------------|---------------|
| <b>Chapter 6</b>     | <b>Accessing LUNs from a Red Hat Enterprise Linux host . . . . .</b>         | <b>43</b>     |
|                      | Loading the HBA driver on the host . . . . .                                 | 45            |
|                      | Configuring the dm-multipath . . . . .                                       | 46            |
|                      | Starting the multipath service. . . . .                                      | 50            |
|                      | Accessing LUNs with dm-multipath support. . . . .                            | 51            |
|                      | Verifying the HBA setting and dm-multipath configuration. . . . .            | 52            |
|                      | Identifying and viewing a list of LUNs . . . . .                             | 55            |
|                      | Creating and mounting a file system on a multipath device . . . . .          | 58            |
|                      | Tuning the dm-multipath configuration . . . . .                              | 60            |
|                      | Stopping the dm-multipath . . . . .                                          | 61            |
|                      | Accessing LUNs without dm-multipath support . . . . .                        | 62            |
|                      | Identifying and viewing LUN information. . . . .                             | 63            |
|                      | Creating a partition . . . . .                                               | 64            |
|                      | Creating a file system . . . . .                                             | 66            |
|                      | Labeling and mounting the file system. . . . .                               | 67            |
|                      | Stopping or reloading the HBA driver and shutting down . . . . .             | 69            |
| <br><b>Chapter 7</b> | <br><b>Accessing LUNs from a SUSE Linux Enterprise Server host . . . . .</b> | <br><b>71</b> |
|                      | Loading the HBA driver on the host . . . . .                                 | 72            |
|                      | Configuring the dm-multipath . . . . .                                       | 73            |
|                      | Starting the multipath service. . . . .                                      | 77            |
|                      | Accessing LUNs with dm-multipath support. . . . .                            | 78            |
|                      | Verifying the HBA setting, dm-multipath configuration . . . . .              | 79            |
|                      | Identifying and viewing a list of LUNs . . . . .                             | 82            |
|                      | Creating and mounting a file system on a LUN . . . . .                       | 85            |
|                      | Tuning the dm-multipath configuration . . . . .                              | 87            |
|                      | Stopping dm-multipath . . . . .                                              | 88            |
|                      | Accessing LUNs without dm-multipath support . . . . .                        | 89            |
|                      | Identifying and viewing the LUN information. . . . .                         | 90            |
|                      | Creating a partition . . . . .                                               | 92            |
|                      | Creating a file system . . . . .                                             | 93            |
|                      | Labelling and mounting the file system . . . . .                             | 94            |
|                      | Stopping or reloading the HBA driver and shutting down . . . . .             | 96            |

|                   |                                                                         |            |
|-------------------|-------------------------------------------------------------------------|------------|
| <b>Chapter 8</b>  | <b>Configuring SAN Boot . . . . .</b>                                   | <b>97</b>  |
| <b>Chapter 9</b>  | <b>Veritas Array Support Library and Array Policy Module. . . . .</b>   | <b>103</b> |
|                   | Overview of the Array Support Library and Array Policy Module . . . . . | 104        |
|                   | Installing the ASL and APM . . . . .                                    | 105        |
|                   | Removing the ASL and APM . . . . .                                      | 109        |
|                   | Obtaining multipathing information . . . . .                            | 110        |
|                   | Understanding ASL error message types. . . . .                          | 113        |
| <b>Appendix A</b> | <b>Troubleshooting . . . . .</b>                                        | <b>115</b> |
| <b>Appendix B</b> | <b>Troubleshooting — Veritas Storage Foundation . . . . .</b>           | <b>121</b> |



# Preface

---

## About this guide

This guide describes how to install the FCP Linux® Host Utilities 4.0 for Native OS and Veritas® on a Linux host and how to set up the host to access the logical units of storage identified by numbers (LUNs) on a NetApp® storage system. It also explains how to configure the storage system to work with the host.

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

## Audience

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

## Terminology

This guide uses the following terms:

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

## Command conventions

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

## Formatting conventions

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

| Formatting convention       | Type of information                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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.

---

---

### **Attention**

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

---



## About this chapter

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

## About Linux versions

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

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

---

### Note

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

---

## About host bus adapters

FCP Linux Host Utilities supports both QLogic® and Emulex® brand host bus adapters (HBAs). Where the setup procedure differs based on HBA model, the relevant information is called out in this guide. If not otherwise marked, the information in this guide applies to all supported versions of HBA models.

---

### Note

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

---

## Topics in this chapter

This chapter includes the following topics:

- ◆ “[The setup procedure](#)” on page 2
- ◆ “[Where to find the information you need](#)” on page 5
- ◆ “[Understanding the Host Utilities components and requirements](#)” on page 9
- ◆ “[Understanding host and storage system communication](#)” on page 12

# The setup procedure

---

## Before you begin

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

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

You can find the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/index.shtml](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml)

## Setup procedure

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>Install and configure HBAs and software on the host, including the following:</p> <ul style="list-style-type: none"><li>◆ Installing the HBAs if not already installed</li><li>◆ Removing any existing HBA drivers</li><li>◆ Downloading and installing the supported HBA driver</li><li>◆ Installing the FCP Linux Host Utilities software</li><li>◆ Setting required HBA parameters</li><li>◆ Installing dm-multipath packages</li><li>◆ Recording the worldwide port names (WWPNs) of the HBAs</li><li>◆ Installing Veritas Storage Foundation, if you want to use the Veritas multipathing solution</li></ul> |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|      | <p><b>Note</b> _____</p> <p>For information on installing Veritas Storage Foundation, see the <i>Veritas Storage Foundation™ Installation Guide</i> for Linux that is shipped along with the software.</p> <hr/> <ul style="list-style-type: none"> <li>❖ Install Array Support Library (ASL) and Array Policy Module (APM)</li> </ul> <p>See Chapter 9, “<a href="#">Veritas Array Support Library and Array Policy Module</a>,” on page 103.</p> <p>See Chapter 2, “<a href="#">Installing HBAs and Software on the Host</a>,” on page 15.</p> |
| 2    | <p>Configure the storage system, including the following:</p> <ul style="list-style-type: none"> <li>◆ Licensing and starting the FCP service</li> <li>◆ Creating LUNs, creating an initiator group (igroup) that contains the host’s FCP WWPN, and mapping the LUNs to that group</li> </ul> <p><b>Note</b> _____</p> <p>The LUN 0 must be mapped to the host.</p> <hr/> <p>See Chapter 3, “<a href="#">Configuring the Storage System for FCP</a>,” on page 29.</p>                                                                            |
| 3    | <p><b>Optional:</b> 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 31.</p>                                                                                                                                                                                                                                                                                                                                       |
| 4    | <p>Configure the Fibre Channel switches and cabling, including the following:</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 HBAs to the switches or directly to the storage system after the LUN 0 is mapped to the host</li> </ul> <p>See Chapter 5, “<a href="#">Configuring FCP cabling and switches</a>,” on page 39.</p>                                              |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 5    | <p>Discover and access the LUNs from the host, including the following:</p> <ul style="list-style-type: none"> <li>◆ Loading the HBA driver modules</li> <li>◆ Configuring multipathing using dm-multipath</li> </ul> <p>See Chapter 6, “<a href="#">Accessing LUNs from a Red Hat Enterprise Linux host</a>,” on page 43, or Chapter 7, “<a href="#">Accessing LUNs from a SUSE Linux Enterprise Server host</a>,” on page 71.</p> |

## Where to find the information you need

---

### About storage area network (SAN) documentation

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

### Where to go for host and storage system information

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

| If you need information about...                           | See...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Known issues, system requirements, and last minute updates | The <i>Release Notes</i> included with this Host Utilities, available at <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 NetApp Interoperability Matrix at <a href="http://now.netapp.com/NOW/products/interoperability">http://now.netapp.com/NOW/products/interoperability</a>, or the NetApp SAN Support Matrix at <a href="http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml">http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml</a>.</li><li>◆ The <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>◆ The <i>Best Practices for Reliability</i> at <a href="http://now.netapp.com/NOW/knowledge/docs/docs.cgi">http://now.netapp.com/NOW/knowledge/docs/docs.cgi</a>.</li><li>◆ The <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>                                                                                                                                                                                                  |

| <b>If you need information about...</b>                                                     | <b>See...</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Supported Fibre Channel SAN topologies                                                      | The NetApp Interoperability Matrix at <a href="http://now.netapp.com/NOW/products/interoperability">http://now.netapp.com/NOW/products/interoperability</a> , or the NetApp SAN Support Matrix at <a href="http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml">http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml</a> .                                                                                                                                                                                                                                                                                                                                                                                  |
| Installing and configuring the HBA in your host                                             | Respective HBA vendor documentation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Installing and configuring a supported version of SnapDrive for UNIX                        | The <i>Installation and Administration Guide</i> for that version of SnapDrive for UNIX.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Managing SAN storage on the storage system                                                  | <ul style="list-style-type: none"> <li>◆ The <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>◆ The <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>◆ The FilerView online Help</li> <li>◆ The <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> |
| Installing Veritas Storage Foundation                                                       | The <i>Veritas Storage Foundation™ Installation Guide</i> for Linux that is shipped along with the software.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Release information about the products in the Veritas Storage Foundation Linux product line | The <i>Veritas Storage Foundation™ Release Notes</i> for Linux that is shipped along with the software.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

| <b>If you need information about...</b>        | <b>See...</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Veritas Volume Manager                         | <ul style="list-style-type: none"> <li>◆ The <i>Veritas™ Volume Manager Administrator's Guide</i> for Linux that is shipped along with the software.</li> <li>◆ The <i>Veritas Volume Manager Troubleshooting Guide</i> for Linux that is shipped along with the software.</li> </ul>                                                                                                                                                                                      |
| Veritas File System                            | <ul style="list-style-type: none"> <li>◆ The <i>Veritas File System Administrator's Guide</i> for Linux that is shipped along with the software.</li> <li>◆ The <i>Veritas™ File System Programmer's Reference Guide</i> for Linux that is shipped along with the software.</li> </ul>                                                                                                                                                                                     |
| Veritas Cluster Server                         | <ul style="list-style-type: none"> <li>◆ The <i>Veritas™ Cluster Server Installation Guide</i> for Linux that is shipped along with the software.</li> <li>◆ The <i>Veritas™ Cluster Server User's Guide</i> for Linux that is shipped along with the software.</li> <li>◆ The <i>Veritas™ Cluster Server Release Notes</i> for Linux that is shipped along with the software.</li> </ul>                                                                                  |
| Veritas Storage Foundation for Oracle RAC      | <ul style="list-style-type: none"> <li>◆ The <i>Veritas Storage Foundation™ for Oracle RAC Installation and Configuration Guide</i> for Linux that is shipped along with the software.</li> <li>◆ The <i>Veritas Storage Foundation™ for Oracle RAC Release Notes</i> for Linux that is shipped along with the software.</li> </ul>                                                                                                                                        |
| Veritas Storage Foundation Cluster File System | <ul style="list-style-type: none"> <li>◆ The <i>Veritas Storage Foundation™ Cluster File System Installation Guide</i> for Linux that is shipped along with the software.</li> <li>◆ The <i>Veritas Storage Foundation™ Cluster File System Administrator's Guide</i> for Linux that is shipped along with the software.</li> <li>◆ The <i>Veritas Storage Foundation™ Cluster File System Release Notes</i> for Linux that is shipped along with the software.</li> </ul> |

# Understanding the Host Utilities components and requirements

---

## What FCP Linux Host Utilities is

FCP Linux Host Utilities 4.0 enables a Linux host to connect to your storage systems. FCP Linux Host Utilities 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

FCP Linux Host Utilities 4.0 consists of the following software components:

- ◆ The `sanlun lun show -p`—provides multipath information for both native device-mapper multipath and Veritas Storage Foundation dynamic multipathing solutions
- ◆ The `qla2xxx_lun_rescan` script that dynamically rescans LUNs when using a QLogic HBA.

There are two options for this script— `qla2xxx_lun_rescan add all` and `qla2xxx_lun_rescan update all`.

- ❖ `qla2xxx_lun_rescan add all` - This option can be used to discover new LUNs that are mapped to all the ports of the host without interrupting IO on other LUNs which are already mapped to the host.
- ❖ `qla2xxx_lun_rescan update all` - This option can be used to remove LUNs which are unmapped from the filer to the host without interrupting IO on other LUNs which are already mapped to the host.

---

### Note

For the native stack, use the `add` option for detecting newly added LUNs and `update` option for removing unmapped or deleted LUNs.

---

### Note

For the Veritas stack, in a SAN boot environment use only the `add` option for detecting newly added LUNs and for removing the unmapped or deleted LUNs. In a non SAN boot environment, use the `add` option for detecting newly added LUNs and `update` option for removing unmapped or deleted LUNs.

---

**Note**

You should not add and delete LUN at the same time, as you cannot run both `add` and `update` options together.

- ◆ An updated `mpath_prio_ontap` callout file that should be used instead of the `/sbin/mpath_prio_netapp` file included in the Linux distribution
- ◆ The `install` and `uninstall` command scripts to add and remove the FCP Linux Host Utilities components
- ◆ The following diagnostic scripts for collecting troubleshooting information:
  - ❖ `linux_info`—collects configuration information about your Linux host
  - ❖ `filer_info`—collects information about the storage system
  - ❖ `qllogic_info`—collects information about QLogic® switches installed in the network
  - ❖ `brocade_info`—collects information about Brocade® switches installed in the network
  - ❖ `cisco_info`—collects information about Cisco® switches installed in the network
  - ❖ `mcdata_info`—collects information about McData® switches installed in the network
  - ❖ `libHBAAPI.so`—support library for `sanlun`
  - ❖ `san_version`—displays the version of the Host Utilities
  - ❖ `SHsupport.pm`—support library for `*info` scripts
  - ❖ `Telnet.pm`—support library for `*info` scripts
  - ❖ `FTP.pm`—support library for `*info` scripts
- ◆ The manual (`man`) pages for `sanlun` and the diagnostic scripts
- ◆ This *Setup Guide*
- ◆ The *Release Notes*
- ◆ The Quick command reference guides

## Multipathing

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

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

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

### **Veritas multipathing**

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

### **Data ONTAP requirements**

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

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

See Chapter 4, “[Migrating a storage system to single\\_image mode](#),” on page 31 for more information.

# Understanding host and storage system communication

---

## About host and storage system communication

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

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

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

## Supported configurations

The following configurations are supported:

- ◆ Fabric-attached SAN

Two variations of fabric-attached SANs are supported:

- ❖ A single-host FCP connection from the HBA to the storage system through a single switch

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

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

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

- ◆ Direct-attached

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

### Note

---

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

---

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





# Installing HBAs and Software on the Host

## Configuration overview

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

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

---

### Note

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

---

## Topics in this chapter

This chapter includes the following topics:

- ◆ [“Installing the FCP Linux Host Utilities software”](#) on page 16
- ◆ [“Removing the previous HBA drivers”](#) on page 18
- ◆ [“Verifying or installing the required multipathing packages”](#) on page 20
- ◆ [“Installing HBAs”](#) on page 21
- ◆ [“Downloading and installing the HBA driver and applications”](#) on page 22
- ◆ [“Setting required HBA and driver parameters”](#) on page 24
- ◆ [“Recording the WWPNs of the HBAs”](#) on page 26

## Requirement for root privileges

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

# Installing the FCP Linux Host Utilities software

---

## Verifying the correct version

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

## Copying or downloading the Host Utilities software

Copy the FCP Linux Host Utilities file (netapp\_linux\_host\_utils\_4.0.tar.gz) from the CD or download it from the NOW™ NetApp on the Web™ site at <http://now.netapp.com/NOW/cgi-bin/software> to a working directory on your Linux host.

### Note

In SUSE Linux Enterprise Server 9 SP3 x86\_64, there is no 32 bit compatibility library for libdevmapper available. The 32 bit compatibility package for device-mapper-multipath can be downloaded from Novell's website at <http://support.novell.com/techcenter/psdb/8ae2a27b3b63ac95fd776298026596eb.html>. This package needs to be installed, to ensure that sanlun work on SUSE Linux Enterprise Server 9 SP3 x86\_64 bit.

## Installing the Host Utilities software

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

| Step | Action                                                                                                                                                                                                        |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Remove any previous version of FCP Linux Host Utilities. Go to the directory where the previous version is installed (default is /opt/netapp/santools) and enter the following command:<br><b>./uninstall</b> |
| 2    | Go to the working directory to which you copied the Host Utilities file.                                                                                                                                      |
| 3    | To uncompress the file, enter the following command:<br><b>gunzip netapp_linux_host_utils_4_0.tar.gz</b>                                                                                                      |

| Step | Action                                                                                                                                                                                 |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 4    | To extract the files, enter the following command:<br><b>tar -xvf netapp_linux_host_utils_4_0.tar</b>                                                                                  |
| 5    | Change to the netapp_linux_host_utils_4_0 directory. By default, this is a subdirectory of the working directory in which you extracted the Host Utilities files in the previous step. |
| 6    | Enter the following command:<br><b>./install</b>                                                                                                                                       |

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

## Removing the previous HBA drivers

---

### Removing existing QLogic HBA drivers

You need to remove any existing HBA driver modules. To remove the QLogic driver modules, perform the following steps.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Log in to the Linux host as root.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 2    | Stop the qlremote process if it is running by entering the following command:<br><b>pkill qlremote</b>                                                                                                                                                                                                                                                                                                                                                                                                       |
| 3    | Change to the directory where the driver installer files are located.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 4    | Uninstall the existing QLogic driver on the host. Ensure that the driver is not in use (that is, usage count of qla2xxx is 1, because it is referenced by qla2300/qla2400 alone). <ul style="list-style-type: none"><li>◆ For a SANsurfer package that is currently installed on the host, run the following command:<br/><b>./qlinstall --uninstall</b></li><li>◆ For a freshly installed host with the bundled drivers loaded, run the following command:<br/><b>modprobe -r qla2300/qla2400</b></li></ul> |

### Removing existing Emulex HBA drivers

You need to remove any existing HBA driver modules. To remove the Emulex driver modules, perform the following steps.

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                  |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | <p>Uninstall the existing Emulex driver on the host. Ensure that the driver is not in use (that is, usage count of lpfc is 0).</p> <ul style="list-style-type: none"> <li>◆ For a freshly installed host with the bundled drivers loaded, run the following command:<br/><b>modprobe -r lpfc</b></li> <li>◆ Run the following command:<br/><b>./lpfc-install --uninstall</b></li> </ul> |

# Verifying or installing the required multipathing packages

---

## Multipathing packages

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

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

# Installing HBAs

---

## **Verify HBA model is supported**

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

## **Do not mix HBA models**

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

## **Hardware installation**

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

# 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. For details about the supported driver versions, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/index.shtml](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml).

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

## Download from vendor's Web site

You must download the HBA driver software package 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.

## Install the driver software

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

A reboot is required after installing the HBA driver package.

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

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

## Installing HBA management software

You should install the HBA vendor's command-line interface (CLI) or graphical user interface (GUI) management packages. The examples in this guide assume you have installed the SANsurfer® FC HBA CLI software, `scli` (for QLogic) or the HBAnyware CLI Client software (Emulex). You can also perform the same tasks using the equivalent GUI if you prefer that.

**QLogic:** The QLogic management CLI (that is, the `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 the multipathing and storage system failover work correctly.

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

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

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

**Note**

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

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

**For Emulex HBAs**

No special settings are required for supported Emulex HBAs. Use the factory default settings.

# Recording the WWPNs of the HBAs

---

## About WWPNs

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

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

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

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

## Obtaining the WWPn for QLogic HBAs

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

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

For more information about using the SANsurfer FC HBA CLI client, see the *SANsurfer FC HBA CLI Application User's Guide*. To download the guide, go 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 at the bottom of most QLogic HBAs.

## Obtaining the WWPN for Emulex HBAs

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

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

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



# Configuring the Storage System for FCP

## Before you begin

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

## Tasks for configuring the storage system

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

- ◆ Ensure that the FCP service is licensed and started.
- ◆ For active/active storage system configurations, verify that the storage system is running a supported cfmode by running the `fcpx show cfmode` command. The support matrix lists the supported cfmode values for your configuration. If you need to change the cfmode, see Chapter 4, "[Migrating a storage system to single\\_image mode](#)," on page 31.

For standby and mixed cfmode, special Fibre Channel switch zoning is required. See "[Zoning with cfmode standby and mixed](#)" on page 40.

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



# Migrating a storage system to single\_image mode

---

## About this chapter

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

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

Using single\_image mode requires dm-multipath on the Linux host.

## Topics in this chapter

This chapter includes the following topics:

- ◆ “[Understanding cfmode](#)” on page 32
- ◆ “[Migrating to single\\_image mode](#)” on page 34

# Understanding cfmode

---

## What cfmode is

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

- ◆ Single System Image (single\_image) cfmode
- ◆ Dual Fabric (dual\_fabric) cfmode
- ◆ Partner cfmode
- ◆ Standby cfmode
- ◆ Mixed cfmode

In this section, only the single\_image cfmode is explained in detail.

## Impact of changing the cfmode setting

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

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

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

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

---

**Note**

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

---

**Prerequisites for changing cfmode**

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

◆ **Save configuration information:**

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

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

---

**Note**

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

---

◆ **Identify affected hosts and initiators:**

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

Enter the following command:

**igroup show**

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

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

# Migrating to single\_image mode

---

## What single\_image mode is

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

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

## Reasons for using single\_image mode

You might want to migrate your configuration to single\_image mode for the following reasons:

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

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

## Resolving LUN mapping conflicts

In single\_image mode, make sure the following mapping conflicts do not occur:

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

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

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

## Tasks for migrating to `single_image` mode

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

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

## Preparing the host

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

| Step | Action                                                                                                             |
|------|--------------------------------------------------------------------------------------------------------------------|
| 1    | Quiesce host I/O and stop any applications accessing LUNs on the storage system.                                   |
| 2    | Disconnect the host by unmounting file systems on LUNs, and then unloading the drivers for the Fibre Channel HBAs. |
| 3    | Proceed to “ <a href="#">Migrating to single_image mode</a> ” on page 36.                                          |

## Migrating to single\_image mode

To migrate the storage system to single\_image mode, complete the following steps.

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Open a window to the console of each storage system controller.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 2    | <p>Log in at the console of the systems you are upgrading, and enter the following command to identify and resolve LUN and igroup mapping conflicts:</p> <pre><b>lun config_check -S</b></pre> <p><b>Result:</b> Mappings that are displayed using the asterisk character (*) indicate an error, and also provide LUN and igroup information that you can use to resolve conflicts. Use the <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 operating system types, 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 the following command:</p> <pre><b>fcp stop</b></pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 5    | <p>Set the cfmode to single_image. On each system console, enter the following command:</p> <pre><b>fcp 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>fcp 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>fcp 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 37.</p> |
| 9    | <p>Return to administrative privileges. On each storage system, enter the following command:</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 WWNN. (One controller in the active/active configuration takes the WWNN of the other.)

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

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

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

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

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

<http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/index.shtml](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml).

## **Reconfiguring the hosts**

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

# Configuring FCP cabling and switches

---

## **Fabric or direct connections**

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

## **Before connecting the Linux host to Fibre Channel switches**

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

## **Topics in this chapter**

This chapter includes the following topics:

- ◆ [“Switched fabric configurations”](#) on page 40
- ◆ [“Direct-connected configurations”](#) on page 42

# Switched fabric configurations

---

## High availability

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

## Configuration and zoning details

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

## Use default HBA and target topology

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

## Single initiator zoning

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

## Zoning with cfmode standby and mixed

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

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

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

**Connect cables**

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

## Direct-connected configurations

---

### High availability

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

### Use default HBA and target topology

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

### Single\_image cfmode required

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

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

```
fcpx show cfmode
```

For information on changing to single\_image cfmode, see Chapter 4, “[Migrating a storage system to single\\_image mode](#),” on page 31.

### Connect cables

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

## About this chapter

This chapter applies only to Red Hat Enterprise Linux hosts. For SUSE Linux Enterprise Server hosts, see Chapter 7, “[Accessing LUNs from a SUSE Linux Enterprise Server host](#),” on page 71.

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

---

### Note

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

---

## Veritas Storage Foundation

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

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

To do this, run the following command:

```
vxdmptune dmp_restore_interval 60
```

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

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

```
vxdmpadm stop restored
```

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

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

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

```
vxdmpadm stat restored
```

## Topics in this chapter

This chapter includes the following topics:

- ◆ “[Loading the HBA driver on the host](#)” on page 45
- ◆ “[Configuring the dm-multipath](#)” on page 46
- ◆ “[Configuring the multipath service to start automatically after reboot](#)” on page 50
- ◆ “[Accessing LUNs with dm-multipath support](#)” on page 51
- ◆ “[Accessing LUNs without dm-multipath support](#)” on page 62
- ◆ “[Stopping the dm-multipath](#)” on page 61
- ◆ “[Stopping or reloading the HBA driver and shutting down](#)” on page 69

# Loading the HBA driver on the host

---

Loading the HBA driver

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

Before you load the driver

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

Starting FCP

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

Be sure you have configured multipathing before starting FCP.

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

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

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

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

# Configuring the dm-multipath

**Note** \_\_\_\_\_  
Skip this section if you are not configuring dm-multipath support or if you are using Verias Storage Foundation.

## Editing the configuration file

To configure the dm-multipath, perform the following steps.

| Step | Action                                        |                                                                                                                                                                                                                                                                                                                            |
|------|-----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Check if the /etc/multipath.conf file exists. |                                                                                                                                                                                                                                                                                                                            |
|      | If...                                         | Then...                                                                                                                                                                                                                                                                                                                    |
|      | The file exists                               | <ul style="list-style-type: none"><li>a. Open the file.</li><li>b. Comment out all the lines.</li><li>c. Append the sample given in “<a href="#">In Red Hat Enterprise Linux 5</a>” on page 48 or “<a href="#">In Red Hat Enterprise Linux 4</a>” on page 48, depending on your version of the operating system.</li></ul> |
|      | The file does not exist                       | Create a new file with the entries as given “ <a href="#">In Red Hat Enterprise Linux 5</a> ” on page 48 or “ <a href="#">In Red Hat Enterprise Linux 4</a> ” on page 48, depending on your version of the operating system.                                                                                               |
| 2    | Save the changes.                             |                                                                                                                                                                                                                                                                                                                            |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | <p>Edit the <code>/etc/multipath.conf</code> file to exclude (blacklist) all devices that do not correspond to any NetApp LUN and which are mapped to your Linux host (that is, blacklist those devices that are not displayed in the <code>sanlun lun show</code> command). Complete the following steps to blacklist the devices.</p> <ol style="list-style-type: none"> <li>Blacklist the non NetApp SCSI devices installed on your host, using the WWID method. To obtain the WWID of a device, run the <code>scsi_id</code> command on the respective device. <p><b>Example:</b> Suppose <code>/dev/sda</code> is a local SCSI drive. To obtain its WWID, run the following command:</p> <pre>scsi_id -gus /block/sda</pre> <p>The output looks like as given below:</p> <pre>SIBM-ESXSMAW3073NC_FDAR9P66067WJ</pre> <p>Add this entry for blacklisting.</p> </li> <li>Blacklist other devices using the devnode method by adding the following: <pre>devnode "^hd[a-z]"</pre> <pre>devnode "^(ram raw loop fd md dm- sr scd st)[0-9]*"</pre> <pre>devnode "^cciss!c[0-9]d[0-9]*"</pre> <p><b>Example:</b> On a Red Hat Enterprise Linux 4 host:</p> <pre>devnode_blacklist { devnode "^(ram raw loop fd md dm-ls scd st)[0-9]*" devnode "^hd[a-z][0-9]*" devnode "^cciss!c[0-9]d[0-9]*p[0-9]*" wwid SIBM-ESXSMAW3073NC_FDAR9P66067WJ }</pre> </li> </ol> |

| Step | Action            |
|------|-------------------|
| 4    | Save the changes. |

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

#### In Red Hat Enterprise Linux 5:

```
defaults
{
    user_friendly_names no
    rr_min_io          128
}
blacklist
{
    wwid SIBM-ESXSMW3073NC_FDAR9P66067WJ
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9] *"
    devnode "^hd[a-z] "
    devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*] "
}
devices
{
    device
    {
        vendor          "NETAPP"
        product          "LUN"
        path_grouping_policy group_by_prio
        getuid_callout    "/sbin/scsi_id -g -u -s /block/%n"
        prio_callout      "/sbin/mpath_prio_ontap /dev/%n"
        features          "1 queue_if_no_path"
        path_checker       readsector0
        failback           immediate
    }
}
```

#### In Red Hat Enterprise Linux 4:

```
defaults
{
    user_friendly_names no
    rr_min_io          128
}
devnode_blacklist
{
    wwid          SIBM-ESXSMW3073NC_FDAR9P66067WJ
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9] *"
```

```

devnode "^hd[a-z] "
devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*] "
}
devices
{
device
{
    vendor            "NETAPP"
    product            "LUN"
    path_grouping_policy group_by_prio
    getuid_callout     "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout       "/sbin/mpath_prio_ontap /dev/%n"
    features           "1 queue_if_no_path"
    path_checker        readsector0
    failback            immediate
}
}

```

For more information about tuning the multipath configuration, see “[Tuning the dm-multipath configuration](#)” on page 60.

## Starting the multipath service

---

### Note

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

---

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

Enter the following command:

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

### Configuring the multipath service to start automatically after

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

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

```
chkconfig --add multipathd
```

```
chkconfig multipathd on
```

## Accessing LUNs with dm-multipath support

---

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

1. Verify the HBA setting and dm-multipath configuration. For more information, see “[Verifying the HBA setting and dm-multipath configuration](#)” on page 52.
2. Identify and view the list of LUNs. For more information, see “[Identifying and viewing a list of LUNs](#)” on page 55.
3. Create and mount the file system. For more information, see “[Creating and mounting a file system on a multipath device](#)” on page 58.
4. Tune the multipath configuration if required. For more information, see “[Tuning the dm-multipath configuration](#)” on page 60.

## Verifying the HBA setting and dm-multipath configuration

---

### Overview

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

### Verifying the HBA settings

#### QLogic:

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

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

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

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

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

Run the -c all option of the scli command:

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

#### Example:

The following is a sample output of the scli command:

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

```

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

```

### **Emulex:**

The default settings work fine in the case of Emulex.

## **Verifying the configuration**

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

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

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

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

Verify that the `multipath -d` is running by issuing the following command:

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

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

```
multipath -ll
```

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

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

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

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

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

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

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

## Identifying and viewing a list of LUNs

---

### Persistently identifying LUNs

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

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

### Viewing list of LUNs

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

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

```
sanlun lun show all
```

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

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

To display multipath information, for both native device-mapper multipath and Veritas Storage Foundation dynamic multipathing solutions, enter the following command:

```
sanlun lun show -p
```

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

```
[root@lnx199-110 ~]# sanlun lun show -p
f3020-92:/vol/vol1/lun2 (LUN 1)          Lun state: GOOD
Lun Size:      3.6g (3879731200)    Filer_CF_State: Cluster Enabled
Protocol: FCP          Filer Partner: f3020-93
DM-MP DevName: mpath3    (360a9800043346852563444717a513571)    dm-1
Multipath-provider: NATIVE
```

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

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

```
[root@lnx199-115 ~]# sanlun lun show -p
f3020-92:/vol/vol1/115-lun2 (LUN 1)          Lun state: GOOD
Lun Size:      2g (2147483648)    Filer_CF_State: Cluster Enabled
Protocol: FCP          Filer Partner: f3020-93
DMP Nodename: FAS30201_1
Multipath-provider: VERITAS
```

| Veritas    | sanlun | Filer     |       |       | Primary | Partner | Veritas     |
|------------|--------|-----------|-------|-------|---------|---------|-------------|
| path       | path   | Path      | /dev/ | Host  | Filer   | Filer   | IO          |
| state      | state  | type      | node  | HBA   | port    | port    | Policy      |
| ENABLED(A) | UP     | PRIMARY   | sdd   | host0 | 1b      | --      | Round-Robin |
| ENABLED    | UP     | SECONDARY | sdb   | host0 | --      | 1a      | Round-Robin |
| ENABLED    | UP     | SECONDARY | sdh   | host1 | --      | 1b      | Round-Robin |
| ENABLED(A) | UP     | PRIMARY   | sdf   | host1 | 1a      | --      | Round-Robin |

## Dynamically rescanning LUNs

To dynamically rescan the LUNs, use the `qla2xxx_lun_rescan` script.

- ◆ The `qla2xxx_lun_rescan` script dynamically rescans LUNs when using a QLogic HBA.

There are two options for this script— `qla2xxx_lun_rescan add all` and `qla2xxx_lun_rescan update all`.

- ❖ `qla2xxx_lun_rescan add all` - This option can be used to discover new LUNs that are mapped to all the ports of the host without interrupting IO on other LUNs which are already mapped to the host.

- ❖ `qla2xxx_lun_rescan update all` - This option can be used to remove LUNs which are unmapped from the filer to the host without interrupting IO on other LUNs which are already mapped to the host.

---

**Note**

For the native stack, use the `add` option for detecting newly added LUNs and `update` option for removing unmapped or deleted LUNs.

---

---

**Note**

For the Veritas stack, in a SAN boot environment use only the `add` option for detecting newly added LUNs and for removing the unmapped or deleted LUNs. In a non SAN boot environment, use the `add` option for detecting newly added LUNs and `update` option for removing unmapped or deleted LUNs.

---

---

**Note**

You should not add and delete LUN at the same time, as you cannot run both `add` and `update` options together.

---

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

# Creating and mounting a file system on a multipath device

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

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

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

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

type is the file system type, such as ext3 .

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

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

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

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

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

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | <p>Add an entry to <code>/etc/fstab</code> as follows, separating each entry with Tab (except the final two zeros, which are separated by a space):</p> <pre><b>device</b>  <b>mount_point</b>  <b>type</b>  <b>defaults</b>  0 0</pre> <p><i>device</i> is the name of the device in the <code>/dev/mpath/</code> directory.</p> <p><i>mount_point</i> is the mountpoint you created for the file system.</p> <p><i>type</i> is the file system type, such as <code>ext2</code> or <code>ext3</code>.</p> <p><b>Example:</b></p> <pre>/dev/mapper/360a980004334686267344374474d4f75 /mnt/ssl/lun_0 ext3 defaults 0 0</pre> |
| 4    | <p>To mount the file system, enter the following command:</p> <pre><b>mount</b> <b>mountpoint</b></pre> <p><b>Example:</b></p> <pre>mount /mnt/ssl/lun_0</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 5    | <p>Do either of the following:</p> <ul style="list-style-type: none"> <li>◆ Enter <b>df</b> to display the mounted file systems. Verify the new file system is listed.</li> <li>◆ Change to the mountpoint directory and try writing and reading a file to verify the file system mounted successfully.</li> </ul>                                                                                                                                                                                                                                                                                                          |

## Tuning the dm-multipath configuration

---

### Tuning multipath parameters

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

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

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

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

The rr\_min\_io 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
}
```

See “[Configuring the dm-multipath](#)” on page 46 for details on blacklisting the local hard drives and other resources that should not be included in the multipathing configuration.

## Stopping the dm-multipath

---

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

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

| Step | Action                                                                                                          |
|------|-----------------------------------------------------------------------------------------------------------------|
| 1    | Unmount all file systems on the FCP LUNs.                                                                       |
| 2    | Flush the multipath maps by entering the following command:<br><b><code>multipath -F</code></b>                 |
| 3    | Stop the multipath daemon by entering the following command:<br><b><code>/etc/init.d/multipathd stop</code></b> |

# Accessing LUNs without dm-multipath support

---

## Overview

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

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

If you are using a switched Fibre Channel fabric, you must zone your switch to prevent the host from having multiple paths to any one LUN. See “[Zoning with cfmode standby and mixed](#)” on page 40. 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 24.

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

## Tasks for accessing LUNs

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

1. Identify and view the list of LUNs. For more information, see “[Identifying and viewing LUN information](#)” on page 63.
2. Create one or more partitions of the LUN. A file system needs a partition, even if the one partition uses all of the space in the LUN. For more information, see “[Creating a partition](#)” on page 64.
3. Create a file system on the partition. For more information, see “[Creating a file system](#)” on page 66.
4. Create a label for the file system and add an entry for the file system’s label in /etc/fstab. Mount the file system. For more information, see “[Labeling and mounting the file system](#)” on page 67.

Accessing LUNs without dm-multipath

# Identifying and viewing LUN information

---

**Persistently  
identifying the LUN**

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

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

**Viewing information  
about LUNs**

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

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

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

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <div>Enter the following command at the host’s shell prompt:</div> <div># <code>sanlun lun show all</code></div> <div><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.</div> <div><b>Example:</b> The following lines show sample <code>sanlun</code> command output:</div> <div># <code>sanlun lun show all</code></div> <div>filer: lun-pathname device filename adapter protocol lun size lun state</div> <div>filer1: /vol/vol_1/lun_1 /dev/sdb1 host0 FCP 4g (4294967296) GOOD</div> <div>filer1: /vol/vol_1/lun_2 /dev/sdc1 host0 FCP 4g (4294967296) GOOD</div> |

# Creating a partition

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

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

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

**sanlun lun show all**

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

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

| Step | Action                                                                                                                                                                                                                                                                                                                             |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Enter <b>/sbin/fdisk device</b> .<br><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    | To accept the default cylinder values, press Enter twice.                                                                                                                                                                                                                                                                          |
| 6    | Enter <b>p</b> to print the partition information.<br><br>Disk /dev/sdb: 64 heads, <b>32 sectors</b> , 10240 cylinders<br>Units = cylinders of 2048 * 512 bytes<br><br>Device Boot      Start              End          Blocks      Id      System<br>/dev/sdb1              1              10240      10485744      83      Linux |

| Step | Action                                                                            |                                                                                                                                                                                                                                                                                                                                                  |
|------|-----------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 7    | <b>If...</b>                                                                      | <b>Then...</b>                                                                                                                                                                                                                                                                                                                                   |
|      | The <code>sectors</code> value from <a href="#">Step 6</a> is not a multiple of 8 | <p>You must set the sectors manually to ensure best performance. Enter <b>q</b> to quit without creating the new partition.</p> <p>For the latest information and specific <code>fdisk</code> parameters, see 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>.</p> |
|      | The <code>sectors</code> value from <a href="#">Step 6</a> is a multiple of 8     | Enter <b>w</b> to save the changes and exit <code>fdisk</code> .                                                                                                                                                                                                                                                                                 |

## Creating a file system

---

### About file systems

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

### Using mkfs to create a file system

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

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

Accessing LUNs without dm-multipath

# Labeling and mounting the file system

---

About labels

A file system label creates a persistent mapping between a LUN and a file system. See “[Persistently identifying LUNs](#)” on page 55 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 mountpoint for the file system.</p> <p><b>Example:</b> <code>mkdir /mnt/filer0_lun0</code></p> |
| 2    | <p>Open the file <code>/etc/fstab</code> 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><code>LABEL=<i>label</i> <i>mount_point</i> auto defaults 0 0</code></pre> <p><i>label</i> is the file system label you created.</p> <p><i>mount_point</i> is the mountpoint you created for the file system.</p> <p><b>Example:</b></p> <pre><code>LABEL=lun0_label /mnt/filer0_lun0 auto defaults 0 0</code></pre> |
| 4    | Save your changes to /etc/fstab and exit the editor.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 5    | Enter <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>Do either of the following:</p> <ul style="list-style-type: none"> <li>◆ Enter <b>df</b> to display the mounted file systems. Verify the new file system is listed.</li> <li>◆ Change to the mountpoint directory and try writing and reading a file to verify the file system mounted successfully.</li> </ul>                                                                                                                                                     |

# 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, offline the LVM volumes if any, flush the multipath maps, and stop the dm-multipath daemon.

| Step | Action                                                                                                                                                                                                                                                                             |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <p>After stopping multipathing, unload the HBA driver.</p> <p><b>QLogic driver:</b> Enter the appropriate command for your HBA model:</p> <pre>modprobe -r qla2400 modprobe -r qla2300</pre> <p><b>Emulex driver:</b> Enter the following command:</p> <pre>modprobe -r lpfc</pre> |

This is to ensure that the HBA driver is not in use, before unloading it.

## Reloading the HBA driver

**For native stack:** To reload the HBA driver, stop the HBA driver. For more information, see “[Stopping the HBA driver](#)” on page 69. 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)
```

**For Veritas Storage Foundation:** To reload the QLogic driver, complete the following steps.

| Step | Action                                                                                                                                                |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Disable the vxconfigd by using the following command:<br><b>vxctl disable</b><br><b>Note</b> _____<br>You can verify it using the vxctl mode command. |
| 2    | Stop the vxesd, vxrelocate and vxnotify processes.                                                                                                    |
| 3    | Reload the driver.                                                                                                                                    |

**Note**\_\_\_\_\_  
 You need to reload the driver when you map new LUNs from the storage system.  
 The host does not recognize the LUNs unless you reload the driver.

## Shutting down the Linux host

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

**multipath -F**

## About this chapter

This chapter applies only to SUSE Linux Enterprise Server hosts. For Red Hat Enterprise Linux hosts, see Chapter 6, “[Accessing LUNs from a Red Hat Enterprise Linux host](#),” on page 43.

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

---

### Note

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

---

## Topics in this chapter

This chapter includes the following topics:

- ◆ “[Loading the HBA driver on the host](#)” on page 72
- ◆ “[Configuring the dm-multipath](#)” on page 73
- ◆ “[Starting the multipath service](#)” on page 77
- ◆ “[Accessing LUNs with dm-multipath support](#)” on page 78
- ◆ “[Stopping dm-multipath](#)” on page 88
- ◆ “[Accessing LUNs without dm-multipath support](#)” on page 89

# Loading the HBA driver on the host

---

**Before you load the driver**

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

**Before you load the driver**

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

**Starting FCP**

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

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

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

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

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

# Configuring the dm-multipath

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

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

\_\_\_\_\_

## Editing the configuration file

To configure the dm-multipath, perform the following steps.

| Step | Action                                        |                                                                                                                                                                                                                                                       |
|------|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Check if the /etc/multipath.conf file exists. |                                                                                                                                                                                                                                                       |
|      | <b>If...</b>                                  | <b>Then...</b>                                                                                                                                                                                                                                        |
|      | The file exists                               | <div>a. Open the file.</div> <div>b. Comment out all the lines.</div> <div>c. Append the sample given “<a href="#">In SUSE Linux Enterprise Server 10:</a>” on page 75 or , depending on your version of the operating system.</div>                  |
|      | The file does not exist                       | Create a new file with the entries as given in the sample “ <a href="#">In SUSE Linux Enterprise Server 10:</a> ” on page 75 or “ <a href="#">In SUSE Linux Enterprise Server 9:</a> ” on page 75, depending on your version of the operating system. |
| 2    | Save the changes.                             |                                                                                                                                                                                                                                                       |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | <p>Edit the <code>/etc/multipath.conf</code> file to exclude (blacklist) all devices that do not correspond to any NetApp LUN and which are mapped to your Linux host (that is, blacklist those devices that are not displayed in the <code>sanlun lun show</code> command). Complete the following steps to blacklist the devices.</p> <ol style="list-style-type: none"> <li>Blacklist the non NetApp SCSI devices installed on your host, using the WWID method. To obtain the WWID of a device, run the <code>scsi_id</code> command on the respective device. <p><b>Example:</b> Suppose <code>/dev/sda</code> is a local SCSI drive. To obtain its WWID, run the following command:</p> <pre>scsi_id -gus /block/sda</pre> <p>The output looks like as given below:</p> <pre>SIBM-ESXSMAW3073NC_FDAR9P66067WJ</pre> <p>Add this entry for blacklisting.</p> </li> <li>Blacklist other devices using the devnode method by adding the following: <pre>devnode "^hd[a-z]"</pre> <pre>devnode "^(ram raw loop fd md dm- sr scd st)[0-9]*"</pre> <pre>devnode "^cciss!c[0-9]d[0-9]*"</pre> <p><b>Example:</b> On a SUSE Linux Enterprise Server 9 host:</p> <pre>devnode_blacklist { devnode "^(ram raw loop fd md dm-ls scd st)[0-9]*" devnode "^hd[a-z][0-9]*" devnode "^cciss!c[0-9]d[0-9]*p[0-9]*" wwid SIBM-ESXSMAW3073NC_FDAR9P66067WJ }</pre> </li> </ol> |

| Step | Action            |
|------|-------------------|
| 4    | Save the changes. |

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

### In SUSE Linux Enterprise Server 10:

```
defaults
{
    user_friendly_names no
    rr_min_io          128
}
blacklist
{
    wwid SIBM-ESXSMAW3073NC_FDAR9P66067WJ
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z]"
    devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
}
devices
{
    device
    {
        vendor          "NETAPP"
        product         "LUN"
        path_grouping_policy group_by_prio
        getuid_callout   "/sbin/scsi_id -g -u -s /block/%n"
        prio_callout     "/sbin/mpath_prio_ontap /dev/%n"
        features         "1 queue_if_no_path"
        path_checker      readsector0
        failback          immediate
    }
}
```

### In SUSE Linux Enterprise Server 9:

```
defaults
{
    user_friendly_names no
    rr_min_io          128
}
devnode_blacklist
{
    wwid SIBM-ESXSMAW3073NC_FDAR9P66067WJ
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
```

```

devnode "^hd[a-z] "
devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*] "
}
devices
{
device
{
    vendor            "NETAPP"
    product            "LUN"
    path_grouping_policy group_by_prio
    getuid_callout     "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout       "/sbin/mpath_prio_ontap /dev/%n"
    features            "1 queue_if_no_path"
    path_checker        readsector0
    failback            immediate
}
}

```

For more information about tuning the multipath configuration, see “[Tuning the dm-multipath configuration](#)” on page 87.

## Starting the multipath service

---

### Note

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

---

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

| Step | Action                                                                                                          |
|------|-----------------------------------------------------------------------------------------------------------------|
| 1    | Start the multipath service by entering the following command:<br><code>/etc/init.d/boot.multipath start</code> |
| 2    | Start the multipath daemon by entering the following command:<br><code>/etc/init.d/multipathd start</code>      |

### Configuring the multipath service to start automatically after

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

| Step | Action                                                                                                                                                                                                                  |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Add the multipath service to the boot sequence by entering the following commands:<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

---

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

1. Verify the HBA setting and dm-multipath configuration. For more information, see “[Verifying the HBA setting, dm-multipath configuration](#)” on page 79.
2. Identify and view the list of LUNs. For more information, see “[Identifying and viewing a list of LUNs](#)” on page 82.
3. Create and mount the file system. For more information, see “[Creating and mounting a file system on a LUN](#)” on page 85.
4. Tune the dm-multipath configuration if required. For more information, see “[Tuning the dm-multipath configuration](#)” on page 87.

## Verifying the HBA setting, dm-multipath configuration

---

### Overview

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

### Verifying the HBA setting

#### QLogic:

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

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

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

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

---

#### Note

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

---

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

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

```
# scli -c all
```

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

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

```

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

```

### **Emulex:**

The default settings work fine in the case of Emulex.

## **Verifying the configuration**

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

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

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

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

Verify that the `multipath -d` is running by issuing the following command:

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

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

```
multipath -ll
```

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

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

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

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

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

**Example:** The following is a sample output.

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

## Identifying and viewing a list of LUNs

---

### Persistently identifying LUNs

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

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

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

### Viewing information about LUNs

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

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

```
sanlun lun show all
```

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

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

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

```
sanlun lun show -p
```

```
[root@lnx199-110 ~]# sanlun lun show -p
f3020-92:/vol/vol1/lun2 (LUN 1)          Lun state: GOOD
Lun Size:      3.6g (3879731200)    Filer_CF_State: Cluster Enabled
Protocol: FCP          Filer Partner: f3020-93
DM-MP DevName: mpath3    (360a9800043346852563444717a513571)    dm-1
Multipath-provider: NATIVE
```

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

## Dynamically rescanning LUNs

To dynamically rescan the LUNs, use the `qla2xxx_lun_rescan` script.

- ◆ The `qla2xxx_lun_rescan` script dynamically rescans LUNs when using a QLogic HBA.

There are two options for this script— `qla2xxx_lun_rescan add all` and `qla2xxx_lun_rescan update all`.

- ❖ `qla2xxx_lun_rescan add all` - This option can be used to discover new LUNs that are mapped to all the ports of the host without interrupting IO on other LUNs which are already mapped to the host.
- ❖ `qla2xxx_lun_rescan update all` - This option can be used to remove LUNs which are unmapped from the filer to the host without interrupting IO on other LUNs which are already mapped to the host.

### Note

For the native stack, use the `add` option for detecting newly added LUNs and `update` option for removing unmapped or deleted LUNs.

### Note

For the Veritas stack, in a SAN boot environment use only the `add` option for detecting newly added LUNs and for removing the unmapped or deleted LUNs. In a non SAN boot environment, use the `add` option for detecting newly added LUNs and `update` option for removing unmapped or deleted LUNs.

**Note**

---

You should not add and delete LUN at the same time, as you cannot run both `add` and `update` options together.

---

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

Accessing LUNs with dm-multipath support

# Creating and mounting a file system on a LUN

---

Overview

Once a LUN is added to the multipath configuration, you can create a file system.  
  
When creating a file system, you use the multipath device in the /dev/mapper/ directory.

Creating the file system

To create a file system on a LUN, enter the following command on the Linux host console:  
  
**mkfs -t *type* /dev/mapper/*device***  
  
*type* is the file system type, such as ext3 .  
  
*device* is the multipath device name of the LUN in the /dev/mapper/ directory.  
  
**Example:** The following command creates an ext3 file system on a LUN:  
  
mkfs -t ext3 /dev/mapper/360a9800043346536684a337a7553424f

Mounting the file system

Multipath devices in the /dev/mapper/ directory that are listed in /etc/fstab are automatically remounted after a reboot, assuming you have correctly configured the multipath service to start automatically.  
  
To mount the new file system, you create a mountpoint, and add an entry to the /etc/fstab file. To mount the file system, complete the following steps.

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3    | <p>Add an entry to <code>/etc/fstab</code> as follows, separating each entry with a tab (except the final two zeros, which are separated by a space):</p> <pre><b>device</b>  <b>mount_point</b>  <b>type</b>  <b>defaults</b>  0 0</pre> <p><i>device</i> is the name in the <code>/dev/mapper/</code> directory.</p> <p><i>mount_point</i> is the mountpoint you created for the file system.</p> <p><i>type</i> is the file system type, such as <code>ext2</code>, <code>ext3</code>, or <code>reiserfs</code>.</p> <p><b>Example:</b></p> <pre>/dev/mapper/360a9800043346536684a337a7553424f /mnt/ssl/lun_0      ext3      defaults      0 0</pre> |
| 4    | <p>Enter the following command to mount the file system:</p> <pre><b>mount -a</b></pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 5    | <p>Do either of the following:</p> <ul style="list-style-type: none"> <li>◆ Enter <b>df</b> to display the mounted file systems. Verify the new file system is listed.</li> <li>◆ Change to the mountpoint and try writing and reading a file to verify the file system mounted successfully.</li> </ul>                                                                                                                                                                                                                                                                                                                                                |

## Tuning the dm-multipath configuration

---

### Tuning multipath parameters

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

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

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

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

The rr\_min\_io 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
}
```

See the “[Configuring the dm-multipath](#)” on page 73 for details on blacklisting the local hard drives and other resources that should not be included in the multipathing configuration.

## Stopping dm-multipath

---

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

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

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

# Accessing LUNs without dm-multipath support

---

## Overview

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

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

If you are using a switched Fibre Channel fabric, you must zone your switch to prevent the host from having multiple paths to any one LUN. See “[Zoning with cfmode standby and mixed](#)” on page 40. 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 24.

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

## Tasks for accessing LUNs

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

1. Identify and view the list of LUNs. For more information, see “[Identifying and viewing the LUN information](#)” on page 90.
2. Create one or more partitions of the LUN. A file system needs a partition, even if the one partition uses all of the space in the LUN. For more information, see “[Creating a partition](#)” on page 92.
3. Create a file system on the partition. For more information, see “[Creating a file system](#)” on page 93.
4. Create a label for the file system and add an entry for the file system’s label in /etc/fstab. Mount the file system. For more information, see “[Labelling and mounting the file system](#)” on page 94.

## Identifying and viewing the LUN information

---

### Persistently identifying LUNs

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

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

### Viewing information about LUNs

You can view information about LUNs in the following ways:

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

### Using lsscsi command

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

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

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

Using the sanlun command

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

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

# Creating a partition

---

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

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

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

**sanlun lun show all**

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

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

| Step | Action                                                                                                                                                       |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Enter the following command:<br><b>fdisk fcp_device</b><br><i>fcp_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

---

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

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | <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 SUSE Linux Enterprise Server 9, the partition name is the name of the link that maps the LUN on the storage system to a local disk device concatenated with the letter p and the partition number.</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> |

Accessing LUNs without dm-multipath

# Labelling and mounting the file system

---

About labels

A file system label creates a persistent mapping between a LUN and a file system. See “[Persistently identifying LUNs](#)” on page 90 for more information.

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

Using e2label to create a file system label

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

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

Mounting the file system

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

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

| Step | Action                                                                                      |
|------|---------------------------------------------------------------------------------------------|
| 5    | To mount the file systems, enter the following command on the host console:<br><b>mount</b> |

# 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, offline the LVM volumes if any, flush the multipath maps, and stop the dm-multipath daemon.

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
```

This is to ensure that the HBA driver is not in use, before unloading it.

## Reloading the HBA driver

To reload the HBA driver, first complete the preceding 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)
```

## 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**

SAN boot is currently supported with Red Hat Enterprise Linux and Veritas Storage Foundation only. For information on the supported configuration, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/index.shtml](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml).

---

**About this chapter**

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

**Configuring SAN boot on the host**

To configure SAN boot on the host using a NetApp storage systems, perform the following steps.

| Step | Action                                                                                         |
|------|------------------------------------------------------------------------------------------------|
| 1    | On the NetApp storage system, choose a supported cfmode.                                       |
| 2    | Create a LUN on the storage system which will act as the SAN boot LUN.                         |
| 3    | Map the LUN to the host.                                                                       |
| 4    | Ensure that only one non-proxy path is available to the LUN.                                   |
| 5    | Ensure that only one SAN boot LUN is available to the host and that LUN is mapped as LUN ID 0. |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6    | <p>Enable the boot BIOS of the HBA port to which the SAN boot LUN is mapped.</p> <p>For information on how to enable the boot BIOS, see the respective HBA vendor-specific documentation.</p> <p><b>Note</b> _____</p> <p>Enable the spinup delay option for the HBA port. This permits the HBA BIOS to scan for 3 minutes until it finds a boot device.</p>                                                                                                                                                                                                                                        |
| 7    | <p><b>Note</b> _____</p> <p>Before installing the OS see the section on rootability in the <i>Veritas™ Volume Manager Administrator's Guide</i> for Linux that is shipped along with the software for partitioning information.</p> <p><b>Attention</b> _____</p> <p>Red Hat Enterprise Linux 4 Update 4 does not have inbox drivers for 4Gb QLogic cards, therefore use the driver kit provided by <a href="#">QLogic Website</a>.</p> <p>After performing the appropriate changes to the HBA BIOS and ensuring the SAN boot LUN is visible, install the operating system on the SAN boot LUN.</p> |
| 8    | <p>After the operating system installation, reboot the host. As a result, the host will boot from the SAN boot LUN on the storage system through a non-proxy path.</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 9    | <p>Install FCP Linux Host Utilities 4.0. For information on installing the FCP Linux Host Utilities 4.0, see “<a href="#">Installing the FCP Linux Host Utilities software</a>” on page 16.</p>                                                                                                                                                                                                                                                                                                                                                                                                     |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 10   | <p>Install the supported version of the drivers.</p> <p>For more information about the installation, see “<a href="#">Installing HBAs</a>” on page 21.</p> <p><b>QLogic:</b></p> <ol style="list-style-type: none"> <li>Open the modprobe.conf file.</li> <li>Append the <code>options qla2xxx</code> statement with settings in the modprobe.conf file on a single line; only one <code>options qla2xxx</code> statement is allowed: <pre>ql2xsuspendcount=6 MaxRetriesPerPath=10 MaxRetriesPerIo=10 ql2xfailover=0</pre> </li> <li>Save the file and exit.</li> <li>Set the QLogic parameters to the recommended values. For more information, see “<a href="#">QLogic HBA parameters</a>” on page 24.</li> <li>Run the following command to reflect the changes in the initrd image. <pre>qlinstall -up</pre> <p>This will ensure that the initrd image is updated with the new driver version without unloading the driver.</p> </li> <li>Reboot the host so that the new driver parameters come into effect.</li> </ol> <p><b>Emulex:</b></p> <ol style="list-style-type: none"> <li>Run the following command to install the Emulex driver. <pre>lpfc-install</pre> </li> </ol> <p><b>Note</b> _____<br/> The <code>lpfc-install</code> command will only update mkinitrd with the requisite driver version and will not unload the driver.<br/> _____</p> <ol style="list-style-type: none"> <li>Reboot the host to load the requisite driver.</li> </ol> |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 11   | Verify the HBA settings. For more information, see “ <a href="#">Verifying the HBA settings</a> ” on page 52.                                                                                                                                                                                                                                                                                                                                                                  |
| 12   | Install Veritas Storage Foundation and respective ASL and APM.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 13   | <p>Encapsulate the root disk for use in VxVM. Any existing data on the disk is preserved by encapsulation when the disk is placed under VxVM control. To encapsulate the disk run the following command:</p> <pre><b>vxdiskadm</b></pre> <p>For detailed steps, see the section on encapsulating the disk in <i>Veritas™ Volume Manager Administrator's Guide</i> for Linux that is shipped along with the software.</p>                                                       |
| 14   | Reboot the host after encapsulation.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 15   | <p>Verify the encapsulation by issuing the following command:</p> <pre><b>vxprint</b></pre> <p>It will display the rootvol and swapvol under the corresponding disk group.</p>                                                                                                                                                                                                                                                                                                 |
| 16   | <p>To make the <code>vxdatapadm restore</code> daemon interval value persistent in a SAN booted host, complete the following steps:</p> <ol style="list-style-type: none"> <li>Edit the <code>/etc/vx/vxvm-startup</code> file.</li> <li>Add the following entries after the <code>xconfigd_opts=\$opts</code> line: <pre><b>/usr/sbin/vxconfigd -k -m enable</b> <b>/usr/sbin/vxdmpadm stop restore</b> <b>/usr/sbin/vxdmpadm start restore interval=60</b></pre> </li> </ol> |
| 17   | Reboot the host.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 18   | <p data-bbox="494 239 1201 300">Configure the paths to the HBA boot BIOS as primary, secondary, tertiary, and so on, on the boot device.</p> <p data-bbox="494 326 1161 387">For more information, see the respective HBA vendor-specific documentation.</p> <p data-bbox="494 413 1233 508"> <b>Note</b> _____<br/> It is recommended to have non-proxy paths as primary paths in the HBA BIOS. </p> |



## About this chapter

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

---

### Note

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

---

## Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Overview of the Array Support Library and Array Policy Module](#)” on page 104
- ◆ “[Installing the ASL and APM](#)” on page 105
- ◆ “[Removing the ASL and APM](#)” on page 109
- ◆ “[Obtaining multipathing information](#)” on page 110
- ◆ “[Understanding ASL error message types](#)” on page 113

# Overview of the Array Support Library and Array Policy Module

---

## What the Array Support Library is

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

The DDL is a component of VxVM that discovers available enclosure information and the attributes of disks and disk arrays that are connected to a host system. The DDL calls ASL functions during the storage discovery process on the host. The ASL in turn “claims” a device based on vendor and product identifiers. The claim associates the storage array model and product identifiers with the device.

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

---

### Note

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

---

## What the Array Policy Module is

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

# Installing the ASL and APM

## ASL and APM installation overview

Installing the ASL and APM involves the following tasks:

- ◆ Verifying that your configuration meets system requirements.  
See the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at [http://now.netapp.com/NOW/knowledge/docs/san/fcp\\_iscsi\\_config/index.shtml](http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml).
- ◆ Determining the ASL and APM version if it is currently installed
- ◆ Obtaining the ASL and APM
- ◆ Installing a fresh version of the ASL and APM or upgrading an existing ASL and APM version

## Determining the ASL and APM version

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

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

```
[root@lnx162 ~]# vxddladm listversion
```

| LIB_NAME          | ASL_VERSION  | Min. VXVM version |
|-------------------|--------------|-------------------|
| libvxnetapp.so    | vm-5.0-rev-1 | 5.0               |
| libvxhitachi.so   | vm-5.0-rev-1 | 5.0               |
| libvxcscovrts.so  | vm-5.0-rev-1 | 5.0               |
| libvxibmds6k.so   | vm-5.0-rev-1 | 5.0               |
| libvxhds9980.so   | vm-5.0-rev-1 | 5.0               |
| libvxibmds4k.so   | vm-5.0-rev-2 | 5.0               |
| libvxveritas.so   | vm-5.0-rev-1 | 5.0               |
| libvxshark.so     | vm-5.0-rev-1 | 5.0               |
| libvxhdsusp.so    | vm-5.0-rev-2 | 5.0               |
| libvxibmds8k.so   | vm-5.0-rev-1 | 5.0               |
| libvxhpalua.so    | vm-5.0-rev-1 | 5.0               |
| libvxxp1281024.so | vm-5.0-rev-1 | 5.0               |
| libvxCLARiON.so   | vm-5.0-rev-1 | 5.0               |
| libvxhdsalua.so   | vm-5.0-rev-1 | 5.0               |
| libvxemc.so       | vm-5.0-rev-2 | 5.0               |
| libvxxp256.so     | vm-5.0-rev-1 | 5.0               |
| libvxhds.so       | vm-5.0-rev-1 | 5.0               |
| libvxpurple.so    | vm-5.0-rev-1 | 5.0               |

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

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

```
[root@lnx162 ~]# vxdmpadm listapm all
```

| Filename       | APM Name    | APM Version | Array Types  | State      |
|----------------|-------------|-------------|--------------|------------|
| dmpap.ko       | dmpap       | 1           | A/P          | Active     |
| dmpap.ko       | dmpap       | 1           | A/P-C        | Active     |
| dmpupalua.ko   | dmpupalua   | 1           | A/A-A-HP     | Not-Active |
| dmpahdsalua.ko | dmpahdsalua | 1           | A/A-A-HDS    | Not-Active |
| dmpjbod.ko     | dmpjbod     | 1           | Disk         | Active     |
| dmpjbod.ko     | dmpjbod     | 1           | APdisk       | Active     |
| dmpaaa.ko      | dmpaaa      | 1           | A/A-A        | Not-Active |
| dmpCLARiion.ko | dmpCLARiion | 1           | CLR-A/P      | Not-Active |
| dmpCLARiion.ko | dmpCLARiion | 1           | CLR-A/PF     | Not-Active |
| dmpaa.ko       | dmpaa       | 1           | A/A          | Active     |
| dmpnetapp.ko   | dmpnetapp   | 1           | A/A-NETAPP   | Active     |
| dmpnetapp.ko   | dmpnetapp   | 1           | A/P-C-NETAPP | Active     |
| dmpnetapp.ko   | dmpnetapp   | 1           | A/P-NETAPP   | Active     |
| dmpapg.ko      | dmpapg      | 1           | A/PG         | Not-Active |
| dmpapg.ko      | dmpapg      | 1           | A/PG-C       | Not-Active |
| dmpapf.ko      | dmpapf      | 1           | A/PF-VERITAS | Not-Active |
| dmpapf.ko      | dmpapf      | 1           | A/PF-T3PLUS  | Not-Active |

## Obtaining the ASL and APM software

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

## Installing ASL and APM software you downloaded

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

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2    | <p>If you already have NetApp storage configured as JBOD in your VxVM configuration, remove the JBOD support for NetApp by entering the following command:</p> <pre><b>vxddladm rmjbod vid=NETAPP</b></pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 3    | <p>Install the ASL and APM according to the directions provided by Symantec. You should have your LUNs set up before you install the ASL and APM.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| 4    | <p>If your host is connected to NetApp storage, you can verify the installation by locating the NetApp Enclosure Type in the output of the <code>vxddmpadm listenclosure all</code> command. The output shows the model name of the storage device if you are using enclosure-based naming with VxVM.</p> <pre><b>vxddmpadm listenclosure all</b></pre> <hr/> <p><b>Note</b></p> <p>After you verify the installation, you can use the <code>vxddctl enable</code> command to restart the <code>vxconfigd</code> daemon. You only use this command when your host is connected to NetApp storage systems.</p> <hr/> <p><b>Example:</b> The <code>vxddmpadm listenclosure all</code> command shows the Enclosure Types as FAS3070 in this example.</p> <pre>[root@lnx162 ~]# vxddmpadm listenclosure all ENCLR_NAME  ENCLR_TYPE  ENCLR_SNO  STATUS      ARRAY_TYPE ===== Disk        Disk        DISKS      CONNECTED   Disk FAS30700    FAS3070     3079851    CONNECTED   A/P-C-NETAPP</pre> |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 5    | <p>If your host is not connected to storage, use the following command:</p> <pre><b>vxddladm listsupport all</b></pre> <p><b>Example:</b> The following is a sample of the output you see when you enter the vxddladm listsupport all command.</p> <pre>[root@lnx162 ~]# vxddladm listsupport all LIBNAME          VID ===== libvxnetapp.so    NETAPP libvxhitachi.so   HITACHI libvxcscovrts.so  CSCOVRTS libvxibmds6k.so   IBM libvxhds9980.so   HITACHI libvxibmds4k.so   IBM libvxveritas.so   VERITAS libvxshark.so     IBM libvxhdsusp.so    HITACHI libvxibmds8k.so   IBM libvxhpalua.so    HP, COMPAQ libvxxpl281024.so HP libvxCLARiion.so  DGC libvxhdsalua.so   HITACHI libvxemc.so       EMC libvxxp256.so     HP libvxhds.so       HITACHI libvxpurple.so    SUN libvxxpl2k.so     HP</pre> |

# Removing the ASL and APM

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

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

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

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1    | Log in to the host as the root user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 2    | Stop all I/O to NetApp LUNs.<br><br><b>Note</b><br>In a Storage Foundation RAC cluster, you must also stop clustering on a node before you remove the ASL and APM.                                                                                                                                                                                                                                                                                                                                                        |
| 3    | <ul style="list-style-type: none"><li>◆ Enter the following command to remove the ASL package:<br/><code>rpm -e asl_rpm_name</code><br/><b>Example:</b> The following command removes a previous version of the ASL.<br/><code>rpm -e VRTSNTAPasl-5.0-1.0_RHEL4</code></li><li>◆ Enter the following command to remove the APM package:<br/><code>rpm -e apm_rpm_name</code><br/><b>Example:</b> The following command removes a previous version of the APM.<br/><code>rpm -e VRTSNTAPapm-5.0-1.0_RHEL4</code></li></ul> |

# Obtaining multipathing information

## What an ASL array type is

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

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

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

## Using VxVM to display available paths

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

| Step        | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |            |       |                |       |        |            |              |   |   |        |            |              |            |     |               |            |              |   |   |        |            |              |   |   |        |            |              |            |     |               |             |              |            |     |               |             |              |            |     |               |     |           |   |   |                |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|-------|----------------|-------|--------|------------|--------------|---|---|--------|------------|--------------|------------|-----|---------------|------------|--------------|---|---|--------|------------|--------------|---|---|--------|------------|--------------|------------|-----|---------------|-------------|--------------|------------|-----|---------------|-------------|--------------|------------|-----|---------------|-----|-----------|---|---|----------------|
| 1           | <div>Enter the following command to view all devices:<br/><br/><b>vxdisk list</b><br/><br/><b>Result:</b> The VxVM management interface displays the vxdisk device, type, disk, group, and status. It also shows which disks are managed by VxVM.<br/><br/><b>Example:</b> The following is a sample output of the vxdisk list command.<br/>[root@lnx162 ~]# vxdisk list<br/><table><tr><th>DEVICE</th><th>TYPE</th><th>DISK</th><th>GROUP</th><th>STATUS</th></tr><tr><td>FAS30700_0</td><td>auto:cdsdisk</td><td>-</td><td>-</td><td>online</td></tr><tr><td>FAS30700_2</td><td>auto:cdsdisk</td><td>FAS30700_4</td><td>dg2</td><td>online shared</td></tr><tr><td>FAS30700_3</td><td>auto:cdsdisk</td><td>-</td><td>-</td><td>online</td></tr><tr><td>FAS30700_4</td><td>auto:cdsdisk</td><td>-</td><td>-</td><td>online</td></tr><tr><td>FAS30700_5</td><td>auto:cdsdisk</td><td>FAS30700_0</td><td>dg1</td><td>online shared</td></tr><tr><td>FAS30700_10</td><td>auto:cdsdisk</td><td>FAS30700_8</td><td>dg1</td><td>online shared</td></tr><tr><td>FAS30700_11</td><td>auto:cdsdisk</td><td>FAS30700_7</td><td>dg2</td><td>online shared</td></tr><tr><td>sda</td><td>auto:none</td><td>-</td><td>-</td><td>online invalid</td></tr></table></div> | DEVICE     | TYPE  | DISK           | GROUP | STATUS | FAS30700_0 | auto:cdsdisk | - | - | online | FAS30700_2 | auto:cdsdisk | FAS30700_4 | dg2 | online shared | FAS30700_3 | auto:cdsdisk | - | - | online | FAS30700_4 | auto:cdsdisk | - | - | online | FAS30700_5 | auto:cdsdisk | FAS30700_0 | dg1 | online shared | FAS30700_10 | auto:cdsdisk | FAS30700_8 | dg1 | online shared | FAS30700_11 | auto:cdsdisk | FAS30700_7 | dg2 | online shared | sda | auto:none | - | - | online invalid |
| DEVICE      | TYPE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | DISK       | GROUP | STATUS         |       |        |            |              |   |   |        |            |              |            |     |               |            |              |   |   |        |            |              |   |   |        |            |              |            |     |               |             |              |            |     |               |             |              |            |     |               |     |           |   |   |                |
| FAS30700_0  | auto:cdsdisk                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | -          | -     | online         |       |        |            |              |   |   |        |            |              |            |     |               |            |              |   |   |        |            |              |   |   |        |            |              |            |     |               |             |              |            |     |               |             |              |            |     |               |     |           |   |   |                |
| FAS30700_2  | auto:cdsdisk                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | FAS30700_4 | dg2   | online shared  |       |        |            |              |   |   |        |            |              |            |     |               |            |              |   |   |        |            |              |   |   |        |            |              |            |     |               |             |              |            |     |               |             |              |            |     |               |     |           |   |   |                |
| FAS30700_3  | auto:cdsdisk                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | -          | -     | online         |       |        |            |              |   |   |        |            |              |            |     |               |            |              |   |   |        |            |              |   |   |        |            |              |            |     |               |             |              |            |     |               |             |              |            |     |               |     |           |   |   |                |
| FAS30700_4  | auto:cdsdisk                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | -          | -     | online         |       |        |            |              |   |   |        |            |              |            |     |               |            |              |   |   |        |            |              |   |   |        |            |              |            |     |               |             |              |            |     |               |             |              |            |     |               |     |           |   |   |                |
| FAS30700_5  | auto:cdsdisk                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | FAS30700_0 | dg1   | online shared  |       |        |            |              |   |   |        |            |              |            |     |               |            |              |   |   |        |            |              |   |   |        |            |              |            |     |               |             |              |            |     |               |             |              |            |     |               |     |           |   |   |                |
| FAS30700_10 | auto:cdsdisk                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | FAS30700_8 | dg1   | online shared  |       |        |            |              |   |   |        |            |              |            |     |               |            |              |   |   |        |            |              |   |   |        |            |              |            |     |               |             |              |            |     |               |             |              |            |     |               |     |           |   |   |                |
| FAS30700_11 | auto:cdsdisk                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | FAS30700_7 | dg2   | online shared  |       |        |            |              |   |   |        |            |              |            |     |               |            |              |   |   |        |            |              |   |   |        |            |              |            |     |               |             |              |            |     |               |             |              |            |     |               |     |           |   |   |                |
| sda         | auto:none                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | -          | -     | online invalid |       |        |            |              |   |   |        |            |              |            |     |               |            |              |   |   |        |            |              |   |   |        |            |              |            |     |               |             |              |            |     |               |             |              |            |     |               |     |           |   |   |                |

| Step | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2    | <p>On the host console, enter the following command to display path information for the device you want:</p> <pre><b>vxddmpadm getsubpaths dmpnodename=device</b></pre> <p><i>device</i> is the name listed under the output of the <code>vxdisk list</code> command.</p> <p><b>Example:</b> The following is sample output.</p> <pre>[root@lnx162 ~]# vxddmpadm getsubpaths dmpnodename=FAS30700_11 NAME STATE[A]    PATH-TYPE[M]  CTRL-NAME ENCLR-TYPE ENCLR-NAME ATTRS ===== sdab ENABLED(A)  PRIMARY      c3        FAS3070   FAS30700   - sdg  ENABLED(A)  PRIMARY      c2        FAS3070   FAS30700   - sdn  ENABLED     SECONDARY    c2        FAS3070   FAS30700   - sdu  ENABLED     SECONDARY    c3        FAS3070   FAS30700   -</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 3    | <p>To obtain path information for a host HBA, enter the following command:</p> <pre><b>vxddmpadm getsubpaths ctrl=controller_name</b></pre> <p><i>controller_name</i> is the controller displayed under “CTRL-NAME” in the output of the <code>vxddmpadm getsubpaths dmpnodename</code> command.</p> <p><b>Result:</b> The output displays information about the paths to the storage system (whether the path is a primary or secondary path). The output also lists the storage system that the device is mapped to.</p> <p><b>Example:</b> The following is sample output.</p> <pre>[root@lnx162 ~]# vxddmpadm getsubpaths ctrl=c2 NAME  STATE[A]    PATH-TYPE[M]  DMPNODENAME ENCLR-TYPE ENCLR-NAME ATTRS ===== sdb   ENABLED     SECONDARY     FAS30700_5  FAS3070    FAS30700   - sdc   ENABLED     SECONDARY     FAS30700_2  FAS3070    FAS30700   - sdd   ENABLED     SECONDARY     FAS30700_0  FAS3070    FAS30700   - sde   ENABLED     SECONDARY     FAS30700_3  FAS3070    FAS30700   - sdf   ENABLED     SECONDARY     FAS30700_4  FAS3070    FAS30700   - sdg   ENABLED(A)  PRIMARY       FAS30700_11 FAS3070    FAS30700   - sdh   ENABLED(A)  PRIMARY       FAS30700_10 FAS3070    FAS30700   - sdi   ENABLED(A)  PRIMARY       FAS30700_5  FAS3070    FAS30700   - sdj   ENABLED(A)  PRIMARY       FAS30700_2  FAS3070    FAS30700   - sdk   ENABLED(A)  PRIMARY       FAS30700_0  FAS3070    FAS30700   - sdl   ENABLED(A)  PRIMARY       FAS30700_3  FAS3070    FAS30700   - sdm   ENABLED(A)  PRIMARY       FAS30700_4  FAS3070    FAS30700   - sdn   ENABLED     SECONDARY     FAS30700_11 FAS3070    FAS30700   - sdo   ENABLED     SECONDARY     FAS30700_10 FAS3070    FAS30700   -</pre> |

Using sanlun to  
obtain multipathing  
information

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

| Step        | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |           |        |       |         |         |             |         |         |      |      |      |       |      |       |       |    |       |       |      |      |     |      |      |        |       |       |       |       |       |       |       |       |         |    |           |     |       |    |    |             |             |    |         |     |       |    |    |             |             |    |         |     |       |    |    |             |         |    |           |     |       |    |    |             |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|--------|-------|---------|---------|-------------|---------|---------|------|------|------|-------|------|-------|-------|----|-------|-------|------|------|-----|------|------|--------|-------|-------|-------|-------|-------|-------|-------|-------|---------|----|-----------|-----|-------|----|----|-------------|-------------|----|---------|-----|-------|----|----|-------------|-------------|----|---------|-----|-------|----|----|-------------|---------|----|-----------|-----|-------|----|----|-------------|
| 1           | <p>On the host command line, enter the following command:</p> <pre>sanlun lun show -p</pre> <p><b>Result:</b> sanlun displays path information for each LUN.</p> <p><b>Example:</b> The following is sample output of the sanlun lun show -p command:</p> <pre>f3070-199-205:/vol/vol1/l1 (LUN 2) Lun state: GOOD Lun Size: 100m (104857600) Filer_CF_State: Cluster Enabled Protocol: FCP Filer Partner: f3070-199-206 DMP Nodename: FAS30700_0 Multipath-provider: VERITAS</pre> <table><tr><td>Veritas</td><td>sanlun</td><td>Filer</td><td></td><td></td><td>Primary</td><td>Partner</td><td>Veritas</td></tr><tr><td>path</td><td>path</td><td>Path</td><td>/dev/</td><td>Host</td><td>Filer</td><td>Filer</td><td>IO</td></tr><tr><td>state</td><td>state</td><td>type</td><td>node</td><td>HBA</td><td>port</td><td>port</td><td>Policy</td></tr><tr><td>-----</td><td>-----</td><td>-----</td><td>-----</td><td>-----</td><td>-----</td><td>-----</td><td>-----</td></tr><tr><td>ENABLED</td><td>UP</td><td>SECONDARY</td><td>sdv</td><td>host3</td><td>--</td><td>0a</td><td>Round-Robin</td></tr><tr><td>ENABLED (A)</td><td>UP</td><td>PRIMARY</td><td>sdr</td><td>host3</td><td>0a</td><td>--</td><td>Round-Robin</td></tr><tr><td>ENABLED (A)</td><td>UP</td><td>PRIMARY</td><td>sdk</td><td>host2</td><td>0c</td><td>--</td><td>Round-Robin</td></tr><tr><td>ENABLED</td><td>UP</td><td>SECONDARY</td><td>sdd</td><td>host2</td><td>--</td><td>0c</td><td>Round-Robin</td></tr></table> | Veritas   | sanlun | Filer |         |         | Primary     | Partner | Veritas | path | path | Path | /dev/ | Host | Filer | Filer | IO | state | state | type | node | HBA | port | port | Policy | ----- | ----- | ----- | ----- | ----- | ----- | ----- | ----- | ENABLED | UP | SECONDARY | sdv | host3 | -- | 0a | Round-Robin | ENABLED (A) | UP | PRIMARY | sdr | host3 | 0a | -- | Round-Robin | ENABLED (A) | UP | PRIMARY | sdk | host2 | 0c | -- | Round-Robin | ENABLED | UP | SECONDARY | sdd | host2 | -- | 0c | Round-Robin |
| Veritas     | sanlun                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Filer     |        |       | Primary | Partner | Veritas     |         |         |      |      |      |       |      |       |       |    |       |       |      |      |     |      |      |        |       |       |       |       |       |       |       |       |         |    |           |     |       |    |    |             |             |    |         |     |       |    |    |             |             |    |         |     |       |    |    |             |         |    |           |     |       |    |    |             |
| path        | path                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Path      | /dev/  | Host  | Filer   | Filer   | IO          |         |         |      |      |      |       |      |       |       |    |       |       |      |      |     |      |      |        |       |       |       |       |       |       |       |       |         |    |           |     |       |    |    |             |             |    |         |     |       |    |    |             |             |    |         |     |       |    |    |             |         |    |           |     |       |    |    |             |
| state       | state                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | type      | node   | HBA   | port    | port    | Policy      |         |         |      |      |      |       |      |       |       |    |       |       |      |      |     |      |      |        |       |       |       |       |       |       |       |       |         |    |           |     |       |    |    |             |             |    |         |     |       |    |    |             |             |    |         |     |       |    |    |             |         |    |           |     |       |    |    |             |
| -----       | -----                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | -----     | -----  | ----- | -----   | -----   | -----       |         |         |      |      |      |       |      |       |       |    |       |       |      |      |     |      |      |        |       |       |       |       |       |       |       |       |         |    |           |     |       |    |    |             |             |    |         |     |       |    |    |             |             |    |         |     |       |    |    |             |         |    |           |     |       |    |    |             |
| ENABLED     | UP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | SECONDARY | sdv    | host3 | --      | 0a      | Round-Robin |         |         |      |      |      |       |      |       |       |    |       |       |      |      |     |      |      |        |       |       |       |       |       |       |       |       |         |    |           |     |       |    |    |             |             |    |         |     |       |    |    |             |             |    |         |     |       |    |    |             |         |    |           |     |       |    |    |             |
| ENABLED (A) | UP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | PRIMARY   | sdr    | host3 | 0a      | --      | Round-Robin |         |         |      |      |      |       |      |       |       |    |       |       |      |      |     |      |      |        |       |       |       |       |       |       |       |       |         |    |           |     |       |    |    |             |             |    |         |     |       |    |    |             |             |    |         |     |       |    |    |             |         |    |           |     |       |    |    |             |
| ENABLED (A) | UP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | PRIMARY   | sdk    | host2 | 0c      | --      | Round-Robin |         |         |      |      |      |       |      |       |       |    |       |       |      |      |     |      |      |        |       |       |       |       |       |       |       |       |         |    |           |     |       |    |    |             |             |    |         |     |       |    |    |             |             |    |         |     |       |    |    |             |         |    |           |     |       |    |    |             |
| ENABLED     | UP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | SECONDARY | sdd    | host2 | --      | 0c      | Round-Robin |         |         |      |      |      |       |      |       |       |    |       |       |      |      |     |      |      |        |       |       |       |       |       |       |       |       |         |    |           |     |       |    |    |             |             |    |         |     |       |    |    |             |             |    |         |     |       |    |    |             |         |    |           |     |       |    |    |             |

## Understanding ASL error message types

---

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

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



## Troubleshooting

### Note

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

### Blacklisting the devices

To blacklist the devices, complete the following steps.

| Step | Action                                                                                                |
|------|-------------------------------------------------------------------------------------------------------|
| 1    | Identify the entries in the sysfs that are not displayed by the <code>sanlun lun show</code> command. |
| 2    | Blacklist those sd entries.                                                                           |

### Unloading of the QLogic module fails

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

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

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

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

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

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

## Linux FCP hosts with QLogic initiators fail to detect LUNs

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

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

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

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

To do this, perform the following steps:

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

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

To do this, perform the following steps:

- ❖ Run the `/opt/QLogic_Corporation/SANsurferCLI/scli`
- ❖ Select the Configure HBA Settings from the Main Menu.

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

## LUNs are not visible on the host

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

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

| Configuration setting     | What to do                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Driver version            | <p>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:</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>                              |
| Cabling                   | Verify that the cables between the host and the storage system are properly connected.                                                                                                                                                                                                                                                                                          |
| Zoning                    | Verify that the zoning is correct. For information on zoning see “ <a href="#">Switched fabric configurations</a> ” on page 40.                                                                                                                                                                                                                                                 |
| HBA Setting               | Verify the HBA settings. For information on verifying the HBA settings, see “ <a href="#">Verifying the HBA settings</a> ” on page 52.                                                                                                                                                                                                                                          |
| 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> .                                                                                                                                                                                                         |
| Initiator login           | <p>Verify that the initiator is logged in to the storage system by entering the <code>fcv 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> |

| Configuration setting | What to do                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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> <ul style="list-style-type: none"> <li>◆ <code>lun show -m</code> displays all LUNs and the igroups to which they are mapped.</li> <li>◆ <code>lun show -g igroup-name</code> displays the LUNs mapped to a specific igroup.</li> </ul>                                                                                                                                                                                                                                                                                             |
| 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 Interoperability Matrix at <a href="http://now.netapp.com/NOW/products/interoperability">http://now.netapp.com/NOW/products/interoperability</a>, or the NetApp SAN Support Matrix at <a href="http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml">http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml</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 `devnode_blacklist` or the `blacklist` section of the `/etc/multipath.conf` file. Ensure that the entries are correctly specified.

For information on blacklisting devices on a Red Hat Enterprise Linux host, see “[Configuring the dm-multipath](#)” on page 46.

For more information about blacklisting devices on a SUSE Linux Enterprise Server host, see “[Configuring the dm-multipath](#)” on page 73.

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

```
multipath -v3
```

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

### **FCP performance is poor**

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

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

### **Error message from sanlun fcp show adapter**

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

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

Make sure the package installing the library is installed & loaded

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

### **LUNs unavailable after storage system failover or giveback**

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

**What to do:** Verify that you have set the required driver and HBA settings. Verify that you have correctly zoned your Fibre Channel switches.

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

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

```
multipath -v3
```



## Troubleshooting — Veritas Storage Foundation

### Note

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

### Error when stopping the fencing driver

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

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

### Note

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

### Starting the vxconfigd

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

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

The corresponding log file where `vxcmp` daemon will store the debug messages is `/etc/vx/vxconfigd.log`

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

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

## Setting the dmp\_log\_level

To set dmp\_log\_level to 4 (highest debugging level), run the following command:

```
vxddmpadm settune dmp_log_level=4
```